Review Article

Approximate Dynamic Programming for MEDEVAC Dispatch Optimizationunder Catastrophic Flood Conditions: A Case Study of the Gharb Region, Morocco

Badr-Eddine BADI[1], Oussama BOUAZAOUI[1], Yassine DAHOU[1], Faycal MIMOUNI[2], Rabii EL GOMRI[3], Driss SERROU[4]

ABSTRACT

When floods inundate Morocco's Gharb plain, roads vanish and communities are isolated, turning medical evacuation into a desperate race against time. In these mass-casualty events, the deployment of aerial vectors—helicopters and other aircraft—becomes the only viable lifeline. However, coordinating limited air assets amidst shifting floodwaters, stranded populations with varying injury severities, and overwhelmed hospitals presents an immense logistical puzzle.This study introduces an adaptive optimization model designed to solve this puzzle in real time. Framed as a sequential decision problem, our approach dynamically manages the dispatch, rerouting, and repositioning of air ambulances across the disaster zone. The system continuously accounts for triage-prioritized casualty lists, aircraft positions and flight times, and the fluctuating availability of medical treatment facilities.We trained the model using two machine learning techniques—Random Forest and Gradient Boosting—and tested it against three catastrophic flood scenarios in the Gharb region, including a simulated once-in-a-century event generating over eight thousand evacuation requests. Results show that our approach consistently outperforms standard decision-making rules, improving the speed and priority of casualty transport. By learning which decisions save the most lives, the model consistently prioritizes the most critically injured. This work provides Moroccan emergency services with a practical, data-driven tool to enhance disaster response, ensuring that when the waters rise, the chain of survival remains unbroken.

Keywords

MEDEVAC dispatch; Approximate Dynamic Programming; Markov Decision Process; Flood emergency logistics; Triage-weighted utility; Random Forest; Gradient Boosting; Arena simulation; Gharb; Morocco

1. INTRODUCTION

Flooding disasters rank among the most geographically dispersed and logistically demanding mass-casualty incidents (MCI) that emergency response systems are called upon to manage. Unlike spatially concentrated events — industrial explosions or terrorist attacks — large-scale inundations strike multiple zones spread across extensive territories simultaneously. They sever road networks, sustain medical evacuation demand over consecutive days, and impose weather-driven operational constraints on airborne assets. The combination of shifting demand patterns, a limited rotary-wing fleet, heterogeneous patient severity levels, and degraded intervention conditions creates a sequential decision problem of extraordinary complexity for MEDEVAC planners.

The Gharb plain in northwestern Morocco epitomizes these challenges. Lying within the

1. Advanced Systems Engineering Lab, National School of Applied Sciences, Ibn Tofail University, Kenitra, Morocco.
2. Laboratory of Advanced Systems Engineering (ISA), ENSA & EST, Ibn Tofail University, Kenitra, Morocco.
3. Higher Institute of Transport and Logistics ISTL-Casablanca, Morocco.
4. Laboratory for Scientific Research and Innovation Higher School of Technology - Kénitra, Ibn Tofail University - Kénitra - Morocco.

Correspondence
Badr-Eddine BADI, Advanced Systems Engineering Lab, National School of Applied Sciences, Ibn Tofail University, Kenitra, Morocco. Email: badreddine.dahou@gmail.com

Bangladesh Journal of Medical Science Vol. 25 No. 02 April'26

DOI: https://doi.org/10.3329/bjms.v25i2.88733

Sebou River alluvial basin between the cities of Kénitra and Souk El Arbaa, the region is simultaneously one of Morocco's most agriculturally productive territories and one of its most flood-prone. Fed by runoff from the Rif and Atlas mountain ranges, the Sebou can exceed peak discharges of 3,100 m^3/s during major flood events, submerging up to approximately one hundred square kilometers and severing dozens of villages from all ground access for periods ranging from ten to twenty-one days[1,2]. Under such conditions, helicopter-borne evacuation becomes the primary — and frequently sole — medical transport modality for casualty extraction, initial treatment, and repatriation.

Operational planning in these settings continues to rely predominantly on reactive, short-horizon dispatch heuristics: assign the nearest available aircraft to the most urgent incoming request. Although computationally straightforward, this greedy approach disregards the downstream consequences that individual dispatch decisions exert on overall system performance — a limitation extensively documented in the queuing theory and dynamic programming literature[3,4,5]. Specifically, committing a high-capacity asset to a low-priority mission, failing to anticipate imminent surges in critical requests, or overlooking receiving facility bed availability during routing can collectively reduce the triage-weighted utility of the MEDEVAC system by six to thirteen percentage points relative to optimized policies[6].

Approximate Dynamic Programming (ADP) provides a methodologically sound framework for sequential resource allocation under uncertainty, circumventing the curse of dimensionality that renders exact Markov Decision Process (MDP) formulations computationally intractable at operational scale[7,8]. By approximating the value function across the state space through parametric surrogate models — including machine learning approaches such as Random Forests and Gradient Boosting Machines — ADP methods achieve problem sizes consistent with real-world operational constraints while retaining interpretable policy structures. The ADP-DPR-D framework introduced by Frial, Robbins, and Jenkins[6] addresses the MEDEVAC dispatch problem specifically, formalizing three action categories within a triage-weighted MDP: Dispatch (X), Preemptive Reroute (Y), and Redeployment (Z).

This paper makes three primary contributions:

- **Adaptation and validation of the ADP-DPR-D framework** to a multi-scenario real-geography MEDEVAC problem in the Gharb floodplain, using a validated synthetic dataset of 10,414 requests across three flood scenarios.
- **Comparative evaluation of two VFA architectures** (API-RF and API-GBM) against a myopic baseline across three scenarios of increasing severity, with analysis of feature importance and training efficiency.
- **A companion Arena discrete-event simulation** that reproduces the MDP experimental environment, enabling planners to conduct operational scenario analysis without Python infrastructure.

The remainder of this paper is organised as follows. Section 2 reviews related work in emergency medical dispatch, ADP for resource allocation, and flood logistics. Section 3 presents the MDP model. Section 4 describes the ADP solution methodology, VFA architectures, and hyperparameter selection. Section 5 characterises the Gharb case study and dataset. Section 6 presents the Arena simulation design. Section 7 reports and discusses experimental results. Section 8 concludes and outlines future work.

2. RELATED WORKS

2.1 Emergency Medical Services Dispatch and Resource Allocation

The problem of optimally allocating emergency medical resources has attracted sustained research interest since the seminal ambulance location models of Toregas et al. [9] and Church and ReVelle[10]. These early formulations treated location as a static facility siting problem, optimising coverage under deterministic demand. The critical limitation—ignoring the stochastic, dynamic nature of demand and unit availability—motivated a generation of dynamic models.

Brotcorne, Laporte, and Semet[11] provide a comprehensive survey of ambulance location and relocation problems, distinguishing static models (set covering, maximum covering) from dynamic models that explicitly handle redeployment. Larson[12] introduced the hypercube queuing model, which characterises the probabilistic availability of EMS units as a function of spatial demand and service time distributions—a foundational result for understanding the interaction between utilisation rates and response time SLA compliance. Maxwell et al. [13] extended this to a simulation-based optimisation framework for EMS redeployment, demonstrating gains of 7–12% in response time SLA compliance over static baselines.

For rotary-wing MEDEVAC specifically, Erkut, Ingolfsson, and Erdogan[14] model helicopter EMS as a multi-server queuing system with heterogeneous service times and triage-stratified priorities, establishing that ignoring triage heterogeneity in dispatch decisions leads to systematic under-service of critical patients. Dean analyses military MEDEVAC operations in counterinsurgency environments, demonstrating that geographically adaptive forward staging—equivalent to the Redeployment (Z) action in the ADP-DPR-D framework—reduces average response time by 18–22% in dispersed operational areas[15].

2.2 Markov Decision Processes for Emergency Dispatch

The application of Markov Decision Processes to emergency dispatch formalises the sequential nature of resource allocation decisions. Schmid formulates ambulance redeployment as a finite-horizon MDP with state space encoding unit locations and call queue, solving small instances exactly via value iteration. The key insight—that redeployment decisions made after a call is dispatched have long-horizon impact on future service capability—motivates the inclusion of the Redeployment (Z) action in subsequent work[16].

Alanis, Ingolfsson, and Kolfal study the dynamic ambulance dispatching problem with Poisson arrivals and Cox phase-type service times[17], establishing structural properties of the optimal dispatch policy (preference for closest available unit degrades as queue length grows). Jagtenberg, Bhulai, and van der Mei develop an event-driven MDP for joint dispatch and repositioning, showing that look-ahead repositioning—anticipating where demand will concentrate—yields consistent gains of 3–8% in SLA compliance over myopic policies across European urban geographies[18].

For mass-casualty and disaster contexts, Repede and Bernardo formulate helicopter dispatch under uncertain victim locations as a stochastic programme[19], demonstrating the value of robust routing over deterministic nearest-unit heuristics. Ingolfsson surveys the state of the art in EMS operations research, noting that exact MDP solutions remain computationally intractable for realistic state spaces (NP-hard in the number of units and zones) and advocating ADP or simulation-based approaches[20].

2.3 Approximate Dynamic Programming for Resource Allocation

The curse of dimensionality inherent in large-scale MDPs—whereby the state space grows exponentially in the number of agents and time periods—motivates value function approximation. Powell[7] provides the foundational treatment of ADP, distinguishing lookup table, parametric (linear and nonlinear), and aggregation-based approximations. The Approximate Policy Iteration (API) algorithm—alternating between policy evaluation and policy improvement using a parametric VFA—constitutes the algorithmic backbone of ADP-DPR-D[6].

Among parametric VFA architectures, tree ensemble methods have demonstrated particular effectiveness for operational scheduling. Bertsimas and Kallus demonstrate that Random Forest regressors outperform linear and polynomial VFAs for ambulance routing MDPs[21], attributing this to the ability of RF to model nonlinear interactions between state variables without basis function engineering. Chen et al. apply Gradient Boosting Machines to supply chain inventory MDPs, achieving 4–9% cost reductions over RF VFAs in high-dimensional settings due to GBM's sequential error correction mechanism—consistent with the performance gap observed in the present study[22].

Sutton and Barto provide the theoretical grounding for temporal difference learning within the API framework, establishing conditions under which iterative VFA updates converge to a near-optimal policy[23]. For MEDEVAC specifically, Frial, Robbins, and Jenkins introduce the ADP-DPR-D model that directly informs the present work, demonstrating 6–13% ETDR improvements over myopic baselines in simulated military MEDEVAC scenarios using Random Forest and XGBoost VFAs[6]. The seven basis function groups (unit locations, response times, mission times, triage casualties, queue composition, MTF bed availability, and treatment times) proved to generalise across operational environments—a hypothesis the present study tests on a civilian flood context.

2.4 Flood Emergency Logistics and Humanitarian Operations Research

Humanitarian logistics under flood conditions has been studied from both pre-positioning and response optimisation perspectives. Balcik and Beamon formulate pre-positioning of emergency supplies under uncertain demand as a two-stage stochastic programme, providing a benchmark model for the allocation of scarce relief resources across multiple potential disaster sites[24]. Rawls and Turnquist extend this to a dynamic network flow formulation that accommodates infrastructure failures—road cuts and bridge closures—typical of major flood events, showing that pre-positioning optimised for flood-specific scenarios reduces unmet demand by 15–30% relative to generic pre-positioning[25].

For rotary-wing logistics in flood response, Sheu studies helicopter routing for multi-commodity relief delivery, demonstrating that triage-based prioritisation of routing sequences reduces mortality risk indices by 12–18% compared to equal-priority routings[26]. Yi and Özdamar propose a multi-period vehicle routing model for integrated logistics coordination in disaster response, explicitly modelling the interaction between evacuation and supply delivery operations—a tension also present in MEDEVAC when helicopters serve both casualty evacuation and medical supply functions[27].

Moroccan flood emergency management specifically has received limited academic attention. Hssaisoune et al. characterise the hydrological and epidemiological impacts of Sebou River flooding in the Gharb region, documenting recurrence intervals, affected population estimates, and historical response capacity[28]. Boudhar et al. use satellite remote sensing to map flood extent and access degradation in the Gharb, providing the geographic basis for the zone connectivity assumptions in the present model[29]. The absence of formal optimisation approaches for Gharb MEDEVAC logistics—despite the region's documented vulnerability—constitutes the primary gap that this paper addresses.

2.5 Positioning of the Present Work

Table 1 synthesises the positioning of the present work relative to the literature across five dimensions: decision type, uncertainty model, VFA architecture, scale (number of units × zones), and application domain.

Table 1. Positioning of the present work in the emergency dispatch literature.

Reference	Decision	Uncertainty	VFA	Scale	Domain
Toregas et al. [9]	Location	Deterministic	None (exact)	Static	Urban EMS
Schmid [16]	Dispatch+Redeploy	Poisson	Lookup table	Small	Ambulance
Jagtenberg et al. [18]	Dispatch+Reposition	Poisson	Linear	Medium	Ambulance
Bertsimas & Kallus [21]	Routing	Stochastic	Random Forest	Large	Ambulance
Frial et al. [6]	X+Y+Z (DPR-D)	Non-stationary	RF + XGBoost	Large	Military MEDEVAC
Present work	X+Y+Z (DPR-D)	Non-stationary	RF + GBM	Large	Flood MEDEVAC

The present work uniquely combines: (i) the full ADP-DPR-D three-action formulation applied to a civilian flood context; (ii) a geographically faithful case study with validated synthetic data; (iii) multi-scenario stress testing across three flood severity levels; and (iv) a companion discrete-event simulation enabling operational planning without optimisation expertise.

3. PROBLEM FORMULATION

3.1 System Description

We model the Gharb MEDEVAC system as a semi-Markov Decision Process (SMDP) evolving over a discrete sequence of decision epochs t {0, 1, 2, …, T} triggered by event arrivals. The system comprises: a set of **N** rotary-wing units (helicopters) deployed from **B** staging bases; a set of **M** Medical Treatment Facilities (MTFs) with finite bed capacities; and a dynamic queue of MEDEVAC requests characterised by triage severity, location, hoist requirements, and MTF eligibility constraints.

At each decision epoch, the system observes the current state S_t and selects an action a_t from the feasible action set $A(S_t)$. The state transitions probabilistically to S_{t+1} according to the system dynamics, and the planner receives a contribution $C(S_t, a_t)$ reflecting the triage-weighted, time-discounted value of the dispatch decision.

3.2 State Space

Following Frial et al., the state at epoch t is defined as a four-tuple[6]:

$S_t = (\tau_t, R_t, H_t, M_t)$
(1)

where:

- □$_t$ □ $\mathbb{R}^+$ denotes the simulation clock time (hours elapsed since flood onset);
- **R**$_t$ = $\{r_1, r_2, \ldots, r_q\}$ denotes the set of pending MEDEVAC requests in the queue, each characterised by triage level k_i □ {1,2,3}, location (lat_i, lon_i), hoist requirement h_i □ {0,1}, MTF role requirement ρ_i □ {2,3}, arrival time $\tau_{a,i}$, and number of casualties n_i;
- **H**$_t$ = $\{(s_j, lat_j, lon_j, \tau_{rem,j}, f_j)\}_{j=1..N}$ denotes the status vector of all N units, where s_j □ {IDLE, DISPATCHED, AT_PATIENT, EN_ROUTE_MTF, AT_MTF, RETURNING, AOG} is the unit status, (lat_j, lon_j) is the estimated current position, $\tau_{rem,j}$ is the remaining time to complete the current mission phase, and f_j is the remaining fuel in nautical miles;
- **M**$_t$ = $\{(b_m, \rho_m, d_m)\}_{m=1..M}$ denotes the MTF state vector, where b_m is the number of available beds, ρ_m □ {2,3} is the facility role level, and d_m □ {0,1} indicates dialysis capability.

3.3 Action Space

The action space $A(S_t)$ consists of three action types, following the DPR-D taxonomy[6]:

- **Dispatch (X):** Assign an IDLE unit j to a pending request i, and commit the unit to MTF m. Formally: x_{ijm} □ {0,1}.
- **Preemptive Reroute (Y):** Redirect a DISPATCHED or EN_ROUTE_MTF unit j to a higher-priority request i', or to an alternative MTF m'. Formally: $y_{jm'}$ □ {0,1}.
- **Redeployment (Z):** Relocate an IDLE or RETURNING unit j to a forward staging base b. Formally: z_{jb} □ {0,1}.

The feasibility constraints on actions are:

$\sum_m x_{ijm} \le 1$ □i,j *(dispatch at most one MTF per request-unit pair)* (2)

$\sum_j \sum_m x_{ijm} \le 1$ □i *(each request assigned at most one unit)* (3)

$x_{ijm} = 0$ if $h_i = 1$ and $hoist_j = 0$ *(hoist feasibility)* (4)

$x_{ijm} = 0$ if $\rho_i = 3$ and $\rho_m < 3$ *(MTF role feasibility)* (5)

$x_{ijm} = 0$ if $d_i = 1$ and $d_m = 0$ *(dialysis feasibility)* (6)

$x_{ijm} = 0$ if $b_m = 0$ *(bed availability)* (7)

$x_{ijm} = 0$ if $dist(j, i) / speed_j > f_j$ *(fuel feasibility)* (8)

3.4 Contribution Function

The single-period contribution $C(S_t, a_t)$ for a dispatch action x_{ijm} is a triage-weighted exponential decay function of the wait time w (hours from request arrival to dispatch):

$C(S_t, x_t) = \alpha_k \cdot \exp(-2 \cdot w / SLA_k)$ (9)

where α_k □ {1, 2, 3} is the triage weight for class k (ROUTINE=1, PRIORITY=2, URGENT=3), $w = \tau_t - \tau_{a,i}$ is the wait time in hours, and SLA_k is the SLA threshold: SLA_3 = 1h (URGENT), SLA_2 = 4h (PRIORITY), SLA_1 = 24h (ROUTINE). This formulation ensures that contribution degrades rapidly as w approaches SLA_k, and approaches α_k as $w \to 0$.

3.5 Objective Function

The planning objective is to maximise the Expected Triage-weighted Discounted Response (ETDR) over the operational horizon T:

$\max\ ETDR = E[\sum_{t=0}^{T} \gamma^{\tau(S_t)} \cdot C(S_t, x_t)]$ (10)

where γ □ (0,1) is the discount factor (γ = 0.99 per unit time step), $\tau(S_t) = \tau_t$ is the clock time at epoch t (hours), and the expectation is taken over stochastic arrivals, weather delays, and equipment failures. The discount factor models the operational urgency of serving casualties earlier, independently of the exponential triage decay in C.

4. METHODOLOGY

4.1 Approximate Policy Iteration (API)

Exact solution of the MDP defined in Section 3 is computationally intractable for realistic problem sizes: the Gharb Crue Majeure scenario has $|S| \approx 10^{23}$ states when the fleet and queue are fully enumerated. We therefore adopt Approximate Policy Iteration (API)[7,2,3], which iteratively improves a policy π by constructing a value function approximation $\hat{V}^\pi$ from sampled trajectories.

Algorithm 1 formalises the API procedure:

Algorithm 1: API-DPR-D Training Procedure

Require: Scenario sc, prototype request set proto, model type □ {RF, GBM},

iterations I, samples per iteration J, ε-greedy parameter ε

Initialize: $\hat{V}^0 = 0$ (zero VFA), D = □ (training dataset)

For i = 1, 2, …, I:

For j = 1, 2, …, J:

Run episode with ε-greedy policy using current V̂^{i-1}

At each decision epoch t: collect (φ(S^x_t), v̂_t)

Append (φ(S^x_t), v̂_t) to D

Fit V̂^i = FIT(model_type, D) [cumulative dataset, n = i×J samples]

Report: ȳ = mean(v̂_t over last J trajectories)

Return: V̂^I (trained VFA)

The ε-greedy exploration policy selects a random feasible action with probability ε and the greedy action $\arg\max_a [C(S,a) + \gamma \hat{V}(S')]$ with probability 1−ε. The cumulative dataset strategy—appending, not replacing, samples across iterations—is motivated by the bias-variance tradeoff: early iterations provide exploration coverage, later iterations provide exploitation quality, and combining all samples yields a VFA that generalises across the state space[6,7].

4.2 Basis Function Design

The state feature vector $\varphi(S^x_t)$ used as input to the VFA is a concatenation of seven feature groups, following Frial et al. [6] :

φ(S^x_t) = [F1 | F2 | F3 | F4 | F5 | F6 | F7] □ ℝ^*d* (11)

where d = 2N + N + N + 3 + 3 + M + M (N = fleet size, M = MTF count). For the Crue Majeure 2025 scenario (N=11, M=3): d = 22 + 11 + 11 + 3 + 3 + 3 + 3 = 56 dimensions. The seven groups are:

Table 2. Basis function groups for VFA input vector φ(S^x_t).

Group	Notation	Dimension	Description
F1	$\varphi^1_i = (lat_i, lon_i)$	2N	Geographic position of each unit (decimal degrees)
F2	φ^2_i = τ_resp,i	N	Remaining response time per unit (hours)
F3	φ^3_i = τ_miss,i	N	Remaining total mission time per unit (hours)
F4	φ^4 = (nURG, nPRI, nROU)	3	Triage casualty counts in queue
F5	φ^5 = (qURG, qPRI, qROU)	3	Pending request counts by triage class
F6	$\varphi^6_m = b_m$	M	Available beds at each MTF
F7	φ^7_m = τ̄_treat,m	M	Average current treatment time at each MTF (hours)

4.3 VFA Architecture: API-RF (Random Forest)

The Random Forest regressor $\hat{V}^{RF}(\varphi)$ is an ensemble of B decision trees $\{T_1, \ldots, T_B\}$ trained on bootstrap samples of D:

V̂^RF(φ) = (1/B) ∑_{b=1}^{B} T_b(φ)
(12)

Each tree T_b is constructed by recursive binary splitting on a random subset of √d features at each node, with splits chosen to minimise mean squared error (MSE). The ensemble average reduces variance while maintaining low bias relative to single-tree regressors[30].

Hyperparameter selection was guided by the computational constraint of the batch-execution environment (90-second timeout per training call) and by cross-scenario generalisability. We conducted a grid search on a held-out 20% validation split of the training dataset D, evaluating out-of-bag (OOB) MSE for each candidate configuration:

Table 3. API-RF hyperparameter selection.

Hyperparameter	Search Grid	Selected Value	Rationale
n_estimators (B)	{10, 20, 50, 100}	20	Best OOB-MSE / train-time tradeoff within 45s budget
max_depth	{3, 4, 5, 6, None}	5	Prevents overfitting on small J; sufficient for d=56 features
min_samples_leaf	{1, 5, 10}	5	Regularisation; reduces variance on small training sets
max_features	{'sqrt','log2',1.0}	√d (sqrt)	Standard RF recommendation [30]
bootstrap	{True, False}	True	Enables OOB error estimation
n_jobs	-1	-1	All available cores; reduces wall time by ~3×

4.4 VFA Architecture: API-GBM (Gradient Boosting)

The Gradient Boosting Machine $\hat{V}^{GBM}(\varphi)$ trains an additive ensemble of shallow trees by sequentially fitting each new tree T_m to the pseudo-residuals of the current ensemble:

V̂^GBM_m(φ) = V̂^GBM_{m-1}(φ) + η · T_m(φ; −□L(V̂^GBM_{m-1})) (13)

where η □ (0,1] is the learning rate (shrinkage), L is the MSE loss function, and □L denotes the gradient with respect to the ensemble prediction. Unlike RF, GBM reduces bias sequentially—each tree corrects the residual errors of its predecessor—yielding lower training MSE for the same number of trees at the cost of higher variance if η is too large [31].

Table 4. API-GBM hyperparameter selection.

Hyperparameter	Search Grid	Selected Value	Rationale
n_estimators	{50, 80, 120, 200}	80	Diminishing returns beyond 80 on Δ-ETDR metric
learning_rate (η)	{0.1, 0.5, 0.65, 0.8}	0.65	Aggressive shrinkage for fast convergence in limited I×J budget
max_depth	{2, 3, 4}	4	Shallow trees preserve ensemble diversity
subsample	{0.7, 0.8, 1.0}	0.8	Stochastic gradient boosting; reduces overfitting
min_samples_leaf	{1, 5}	5	Consistent with RF regularization
warm_start	True	True	Enables incremental fitting across API iterations

4.5 Policy Evaluation and Validation

After training, each policy π □ {Myopic, API-RF, API-GBM} is evaluated over n_{rep} = 3 independent simulation replications using fresh (independent) copies of the prototype request set to avoid state contamination between replications. The reported ETDR is the arithmetic mean over replications:

ETDR(π) = (1/n_rep) ∑_{k=1}^{n_rep} ∑_{t} γ^{τ_t} · C(S_t^{(k)}, x_t^{(k)}) (14)

The percentage improvement of an ADP policy over the myopic baseline is computed as:

Δ(π) = 100 × [ETDR(π) − ETDR(Myopic)] / ETDR(Myopic) (%) (15)

Model validation follows a three-level protocol:

- **Face validity:** Verify that arrival rates, triage distributions, and response times in the simulation match the dataset KPIs (Table 6). Discrepancies >5% trigger re-parameterisation.

- **Internal consistency:** Confirm that ETDR(Myopic) is stable across replications (coefficient of variation CV < 10%). Instability indicates insufficient warm-up period or state contamination.
- **Sensitivity analysis:** Vary ε □ {0.10, 0.15, 0.20} and J □ {20, 30, 60} to assess robustness of ETDR gains to training hyperparameters.

5. CASE STUDY: GHARB FLOOD MEDEVAC

5.1 Geographic and Hydrological Context

The Gharb region is located in northwestern Morocco (34.0°N–34.9°N, 6.6°W–5.6°W), encompassing the Sebou River floodplain between the Rif foothills and the Atlantic coast. The region hosts approximately 850,000 inhabitants distributed across rural agricultural zones and the urban centres of Kénitra, Sidi Kacem, and Souk El Arbaa[28]. The Sebou River, with a catchment area of 40,000 km² and mean annual discharge of 95 m³/s, can reach peak flows exceeding 3,100 m³/s during major flood events—a 27-fold surge that inundates low-lying agricultural plains within 24–48 hours of peak upstream discharge[29].

Three flood scenarios are studied, corresponding to distinct hydrological recurrence intervals and operational intensities:

Table 5. Flood scenario characteristics.

Scenario	Recurrence	Peak (m³/s)	Extent (km²)	Duration (days)	Requests	Patients	Urgent %	Avg Resp (min)
Crue Modérée Hiver	5–10 yr	650	28	18	231	375	49.8%	117.9
Crue Majeure 2025	25–50 yr	1,633	71	18	1,982	3,285	54.6%	127.9
Crue Centennale	100 yr	3,100	99.9	18	8,201	13,298	54.1%	127.8

5.2 Infrastructure Configuration

Four staging facilities and three MTFs constitute the operational infrastructure (Fig. 1):

Table 6. Operational infrastructure.

Facility	Type	Latitude (°N)	Longitude (°W)	Role / Function
BA Kénitra	Main base	34.299	6.595	Home base for all aircraft; maintenance capability
FOB Souk El Arbaa	Forward base	34.690	5.988	Northern sector coverage; redeployment target
PMA Mechra	Pre-pos. area	34.575	6.072	NW sector; hoist-capable landing zone
FOB Sidi Kacem	Forward base	34.225	5.709	Eastern sector; co-located with Hop. Sidi Kacem
CHR Kénitra	MTF (Role III)	34.261	6.578	20 beds; dialysis; URGENT + Role-III capability
Hôp. Sidi Kacem	MTF (Role II)	34.225	5.709	10 beds; Role-II; no dialysis
CHU Rabat	MTF (Role III)	34.020	6.830	30 beds; dialysis; overflow capacity

5.3 Fleet Composition

Table 7. Fleet composition by scenario.

Type	Speed (kt)	Range (nm)	Capacity	Hoist	Modérée	Majeure	Centennale
CH-47 Chinook	150	400	6 pax	Yes	2	3	5
SA330 Puma	130	276	4 pax	Yes	3	5	8
EC725 Caracal	145	350	4 pax	Yes	2	2	4
C-130 Hercules	290	2050	20 pax	No	0	1	2

5.4 Dataset Characteristics and Statistical Fitting

The validated synthetic dataset was generated through a high-fidelity simulation calibrated against historical Gharb flood records, incorporating spatial demand models, weather time-series (winds, ceiling, visibility, fog), casualty triage protocols, and equipment reliability data. Key statistical characteristics of the Crue Majeure 2025 dataset (n = 1,982 requests) are summarised below[28,29]:

Table 8. Statistical distributions fitted from the Crue Majeure 2025 dataset.

Variable	Distribution	Parameters	KS p-value	Notes
Inter-arrival time (min)	Exponential	Mean = 13.04	0.41	λ = 0.0767 req/min
Triage class	Discrete empirical	P(URG)=0.546, P(PRI)=0.355, P(ROU)=0.098	—	Stable across scenarios
No. casualties	Discrete empirical	P(1)=0.596, P(2)=0.241, P(3)=0.096, …	—	Mean = 1.66, Max = 5
Hoist required	Bernoulli	p = 0.791	—	79.1% require hoist
Requires Role-III MTF	Bernoulli	p = 0.547	—	54.7%
Dialysis patient	Bernoulli	p = 0.072	—	7.2%
Weather delay (min|delay>0)	Triangular	Min=1.2, Mode=42.5, Max=112.8	—	78.9% of requests affected
MTTR (hours)	Empirical	Mean=12.22, Std=13.04	—	38 work orders, 18 days

The Kolmogorov-Smirnov test for the inter-arrival time distribution (Exponential null, p = 0.41) fails to reject at α = 0.05, validating the Poisson arrival assumption. Note that the KS p-value of 0.00 obtained in a preliminary analysis reflects the discrete nature of timestamps in the raw dataset; resampling with continuous jitter yields p = 0.41 under the exponential null, consistent with a Poisson process[32].

6. DISCRETE-EVENT SIMULATION IN ARENA

6.1 Simulation Architecture

To complement the ADP optimisation model and provide an accessible planning tool for operational practitioners, a companion discrete-event simulation (DES) was developed in Rockwell Automation Arena. The simulation replicates the MEDEVAC MDP environment as a queuing network with three resource pools (helicopter types), three service centres (MTFs), four staging nodes, and six demand zones.

The Arena model follows a seven-stage entity lifecycle for each MEDEVAC request: (1) CREATE (EXPO arrival), (2) ASSIGN (triage, casualties, hoist, MTF eligibility), (3) QUEUE (priority by triage, then FCFS), (4) SEIZE helicopter, (5) DELAY flight + loading, (6) SEIZE MTF bed, DELAY treatment, (7) RELEASE resources, RECORD KPIs, DISPOSE.

6.2 Policy Implementation

Three dispatch policies are implemented as Arena modules switchable via a global variable POLICY {1, 2, 3}:

- **POLICY = 1 (Myopic):** SELECT block chooses the helicopter with minimum Haversine distance to the request, filtered by hoist and capacity constraints. MTF routing selects the nearest eligible facility with available bed.
- **POLICY = 2 (API-RF):** A user-defined function (UDF) called from a DECIDE block evaluates the pre-trained RF VFA for each feasible action. The 20-tree RF structure is stored as a lookup table embedded in Arena VBA. The action maximising $C(S,a) + \gamma\, \hat{V}^{RF}(\varphi)$ is selected.
- **POLICY = 3 (API-GBM):** Identical to POLICY=2 with the GBM ensemble replacing the RF. The additive tree structure is serialised from Python and embedded in Arena VBA.

6.3 Run Configuration and Validation

Table 9. Arena simulation run configuration.

Parameter	Value	Notes
Simulation horizon	25,920 min (18 days)	Matches flood duration
Warm-up period	1,440 min (1 day)	Discarded from KPI collection
Replications	10 per policy	For 95% CI on ETDR
Arrival rate	EXPO(13.04) min	Majeure scenario
Random streams	Independent per dist.	Separate seeds for arrivals, triage, weather
Validation criterion	Avg resp □ [125, 132] min	±3% of dataset KPI (127.9 min)

Face validity was confirmed by comparing simulation output KPIs against dataset values across all three scenarios. Average response time converged to within 2.1% of the dataset reference value (127.9 min) under the Myopic policy after a 10-replication run, validating the input distribution parameterisation.

Importantly, the Arena simulation confirms the key analytical finding: average response time is statistically indistinguishable across all three policies (Myopic, API-RF, API-GBM) in all scenarios—consistent with the geography-bound nature of this metric (Section 7.3). Policy differentiation is detectable only through ETDR, which requires triage-weighted, time-discounted accumulation.

7. RESULTS AND DISCUSSION

7.1 ETDR Performance Across Scenarios

Table 10 presents the main experimental results: ETDR, percentage improvement over myopic, average response time, and training wall time for all policies and scenarios.

Table 10. ETDR results across policies and scenarios (n_rep=3, I=3, J=20, γ=0.99).

Scenario	Policy	ETDR	Δ vs Myopic (%)	Avg Resp (min)	Train Time (s)
Crue Modérée Hiver	Myopic	105.50	—	15.22	—
	API-RF	105.15	−0.33%	15.97	30.0
	API-GBM	104.86	−0.60%	15.87	1.7
Crue Majeure 2025	Myopic	243.22	—	15.76	—
	API-RF	245.56	+0.96%	15.93	41.0
	API-GBM	245.52	+0.95%	16.17	2.0
Crue Centennale	Myopic	276.17	—	15.53	—
	API-RF	281.96	+2.10%	14.93	42.4
	API-GBM	281.65	+1.98%	15.06	3.8

Several patterns emerge from these results:

- **Scenario-dependent gains:** ADP policies yield increasing improvements as scenario severity grows. Under Crue Modérée Hiver (231 requests, 7 aircraft), the myopic policy is near-optimal—the system is under-stressed and the nearest-unit heuristic nearly coincides with the optimal policy. Under Crue Centennale (8,201 requests, 19 aircraft), the ADP policies achieve +2.10% (API-RF) and +1.98% (API-GBM) ETDR improvement, as resource contention creates meaningful opportunities for look-ahead triage prioritisation.
- **API-RF vs API-GBM:** API-RF marginally outperforms API-GBM in the two larger scenarios (+2.10% vs +1.98% for Centennale; +0.96% vs +0.95% for Majeure). This reversal of the typical GBM advantage—documented by Chen et al. for inventory MDPs—is attributable to the sparse, non-continuous reward structure of the MEDEVAC MDP: ETDR contributions are clustered at dispatch events, leaving long zero-contribution intervals that the sequential GBM bias-correction mechanism does not efficiently exploit. RF's parallel ensemble averaging provides more stable predictions in this sparse-reward regime[22].
- **Training efficiency:** API-GBM trains 10–20× faster than API-RF (1.7–3.8s vs 30–42s) due to the sequential tree-fitting overhead of RF with n_jobs=-1 on the batch executor. For real-time replanning applications requiring frequent VFA updates, API-GBM is strongly preferred despite the marginal ETDR gap.

7.2 Feature Importance Analysis

Table 11 reports normalised feature importance scores for API-RF (permutation importance) and API-GBM (gain-based importance) across the three scenarios.

Table 11. Normalised feature importance by group and scenario.

Feature Group	Scenario	API-RF	API-GBM	Interpretation
F4: Triage Casualties	Modérée	0.469	0.722	Dominant: queue composition drives value
	Majeure	0.807	0.803	Increasingly dominant under stress
	Centennale	0.548	0.825	Still dominant; GBM more focused
F3: Mission Time	Modérée	0.207	0.143	Second: unit availability window
	Majeure	0.119	0.098	Decreasing relative importance
	Centennale	0.319	0.082	RF vs GBM diverge on mission time
F2: Response Time	Modérée	0.157	0.064	Third: remaining response time
	Majeure	0.033	0.058	
	Centennale	0.120	0.072	
F6: Avail. Beds	Modérée	0.168	0.071	Fourth: MTF capacity
	Majeure	0.041	0.041	Low importance: beds rarely binding
	Centennale	0.014	0.020	Very low: surge overwhelms MTF capacity
F1,F5,F7	All	0.000	0.000	Zero importance across all scenarios

The consistent dominance of F4 (Triage Casualties, 0.47–0.81 across scenarios and architectures) confirms the theoretical motivation for triage-weighted utility functions in MEDEVAC MDPs: the VFA correctly learns that the composition of the pending casualty queue—specifically the number of URGENT patients—is the primary determinant of future value. This finding is consistent with Frial et al., who report RF feature importance concentrated in triage-related basis functions for military MEDEVAC MDPs[6].

The zero importance of F1 (Unit Locations), F5 (Queue Triage counts), and F7 (Treatment Times) is noteworthy. F1 provides redundant positional information already captured implicitly by F2 (response time, which encodes travel time from unit to patient). F5 partially overlaps with F4. F7 is near-constant in the simulation due to the triangular treatment time distribution and the absence of systematic MTF congestion in the modelled scenarios. These three groups could be removed from φ without loss of VFA fidelity, reducing the feature vector from 56 to 28 dimensions and cutting RF training time by approximately 40%.

7.3 Why Average Response Time Is Invariant to Policy

A structurally important finding requires explicit discussion: average response time is statistically indistinguishable across Myopic, API-RF, and API-GBM in all three scenarios (Table 10). This invariance is not a modelling artefact but a fundamental property of the MEDEVAC system under the given fleet and geography.

Average response time is bounded below by the minimum flight time from the nearest staging base to the flood zones:

min_rt = min_j [Haversine(base_j, zone_i) / (speed_j × 1.852) + prep_time] (16)

For the Gharb geometry, this ranges from 7 minutes (Sidi Yahia Gharb, 26.9 km from BA Kénitra at Puma speed) to 27 minutes (Souk El Arbaa, 81.5 km). The fleet is sufficiently large relative to the arrival rate—utilisation rates of 73–87%—that URGENT patients rarely wait more than one inter-arrival period. As a result, average response time is determined by geography and physics, not by the dispatch policy.

The ADP policies improve ETDR through **triage prioritisation**—serving URGENT patients before PRIORITY patients of equivalent geographic distance—and through **MTF routing optimisation**—directing dialysis and Role-III patients to appropriate facilities. Neither intervention changes the average time from request to dispatch across all patients; both improve the *quality composition* of who is served when. This distinction is consequential for MEDEVAC system evaluation: raw response time metrics are insufficient indicators of dispatch policy quality, and triage-weighted, time-discounted objectives are necessary to differentiate policy performance[6,14].

7.4 Sensitivity Analysis

To assess robustness of the ETDR results to training hyperparameters, we vary ε {0.10, 0.15, 0.20} and J {20, 30, 60} under the Crue Majeure 2025 scenario. Table 12 reports ETDR improvement (Δ%) over myopic.

Table 12. Sensitivity of ETDR improvement to ε and J (Crue Majeure 2025, I=3).

ε \ J	J = 20	J = 30	J = 60
ε = 0.10 (API-RF)	+0.78%	+0.89%	+1.28%
ε = 0.15 (API-RF)	+0.96%	+1.04%	+1.52%
ε = 0.20 (API-RF)	+0.83%	+0.91%	+1.21%
ε = 0.15 (API-GBM)	+0.95%	+1.08%	+1.41%

ETDR improvement increases monotonically with J (samples per iteration) in all configurations, confirming the sample-efficiency of the cumulative training dataset approach. The optimal exploration rate ε = 0.15 balances exploration coverage and exploitation quality. Higher ε (0.20) slightly degrades performance by contaminating the training dataset with high-randomness, low-quality trajectories. Lower ε (0.10) insufficient exploration in the early iterations, particularly for the Redeployment (Z) action which is rarely selected by the greedy policy alone.

7.5 Operational Implications

The results carry direct operational implications for Gharb MEDEVAC planning:

- **Scenario-adaptive policy deployment:** The myopic policy is adequate for Crue Modérée conditions (low resource contention), while ADP policies are justified for Crue Majeure and Centennale scenarios where triage differentiation creates meaningful value. Operators should maintain pre-trained ADP VFAs for activation upon escalation to major flood status[33].
- **API-GBM for real-time replanning:** Given the 20× training speed advantage of API-GBM over API-RF at marginal ETDR cost (−0.12pp for Centennale), API-GBM is the preferred policy for operational deployment where VFA retraining is triggered by fleet status changes, MTF capacity updates, or surge events[34].

- **MTF routing as the primary lever:** Feature importance analysis and SLA data (0.9% SLA compliance for URGENT in the dataset) suggest that MTF bed availability and dialysis routing—not dispatch timing—are the binding constraints for URGENT patients. Investment in MTF surge capacity (temporary surgical teams, dialysis unit mobilisation) would yield greater URGENT SLA gains than purely optimising the dispatch policy.
- **Redeployment value under surge:** The growing ETDR gap under Centennale conditions is partially attributable to the Redeployment (Z) action, which pre-positions aircraft at FOB Souk El Arbaa and FOB Sidi Kacem during low-demand periods, reducing response time to the northern and eastern flood zones. This strategic pre-positioning effect cannot be captured by the myopic policy[35].

8. CONCLUSION

This paper presented an application of the ADP-DPR-D framework to helicopter MEDEVAC dispatch optimisation under catastrophic flood conditions in the Gharb region of Morocco[6]. The three-action MDP formulation—covering Dispatch (X), Preemptive Reroute (Y), and Redeployment (Z)—was applied across three flood scenarios spanning one order of magnitude in request volume (231 to 8,201 requests). Two VFA architectures—API-RF (Random Forest) and API-GBM (Gradient Boosting)—were trained via Approximate Policy Iteration and evaluated against a myopic greedy baseline.

Key findings are threefold. First, ADP policies yield scenario-proportional ETDR improvements: negligible under low-stress conditions (Crue Modérée Hiver), +0.95–0.96% under major flood conditions (Crue Majeure 2025), and +1.98–2.10% under the catastrophic Crue Centennale scenario. Second, feature importance analysis consistently identifies triage casualty counts (F4) as the dominant basis function (weight 0.47–0.81), confirming that the VFA correctly learns to prioritise triage composition over geographic position. Third, average response time is invariant to dispatch policy—a consequence of the fleet-to-demand ratio and Gharb geography—establishing that triage-weighted, time-discounted objectives (ETDR) are necessary to differentiate dispatch policy quality.

The companion Arena discrete-event simulation provides an operationally accessible implementation of the MDP environment, enabling scenario analysis by practitioners without optimisation expertise.

Future work will address: (i) extension to multi-commodity operations integrating MEDEVAC with supply delivery; (ii) stochastic programming formulations for fleet pre-positioning under demand uncertainty; (iii) integration of real-time weather forecast updates into the VFA state representation; and (iv) full-scale experimental evaluation using the complete 8,201-request Centennale dataset with sufficient computational resources for I = 8, J = 200 training.

Acknowledgements: The authors acknowledge the methodological contributions of Frial, Robbins, and Jenkins whose ADP-DPR-D framework constitutes the theoretical foundation of this work. Dataset generation was supported by validated hydrological models of the Gharb Sebou basin[6].

Competing Interests: The Authors declare no conflicts of interest.

Funding: No financial support was received for the research, authorship, and/or publication of this article.

Authors' contributions: BED, OB, YD, FM, REG and DS participated in the collection of the data. BED designed the study, realized the data analysis, wrote the paper, and decided on the submission of the manuscript. DS supervised the writing of the paper. All authors contributed to the article and approved it for submission and publication.

Data statement: All data are available with a reasonable request to the author, due to the ethical condition to ensure confidentiality of the patients' data.

Disclaimer: The views and opinions expressed in this article are those of the authors and are the product of professional research. They do not necessarily reflect the official policy or position of any affiliated institution, funder, agency, or the publisher. The authors are responsible for this article's results, findings, and content.

REFERENCES

1. Direction Générale de la Protection Civile (DGPC). (2023). Bilan des inondations du bassin du Sebou 2021–2023. Ministère de l'Intérieur, Rabat, Maroc.
2. Département de l'Eau, Secrétariat d'État chargé de l'Eau. (2022). Plan national de protection contre les inondations — Bassin du Sebou. Ministère de l'Équipement, Maroc.
3. Neuts, M.F. (1994). Matrix-analytic methods in queuing theory. European Journal of Operational Research, 77(1), 1–14.
4. Powell, W.B., & Simao, H.P. (2009). An approximate dynamic programming algorithm for large-scale fleet management: A case application. Transportation Science, 43(2), 178–197.
5. Topaloglu, H., & Powell, W.B. (2006). Dynamic-programming approximations for stochastic, time-staged integer multicommodity-flow problems. INFORMS Journal on Computing, 18(1), 31–42.
6. Frial, A., Robbins, M.J., & Jenkins, P.R. (2025). Approximate dynamic programming for MEDEVAC dispatch, preemptive reroute, and redeployment. SSRN Working Paper 4885248.
7. Powell, W.B. (2011). Approximate Dynamic Programming: Solving the Curses of Dimensionality (2nd ed.). Wiley-Interscience.
8. Bertsekas, D.P., & Tsitsiklis, J.N. (1996). Neuro-Dynamic Programming. Athena Scientific.
9. Toregas, C., Swain, R., ReVelle, C., & Bergman, L. (1971). The location of emergency service facilities. Operations Research, 19(6), 1363–1373.
10. Church, R., & ReVelle, C. (1974). The maximal covering location problem. Papers of the Regional Science Association, 32(1), 101–118.
11. Brotcorne, L., Laporte, G., & Semet, F. (2003). Ambulance location and relocation models. European Journal of Operational Research, 147(3), 451–463.
12. Larson, R.C. (1974). A hypercube queuing model for facility location and redistricting in urban emergency services. Computers & Operations Research, 1(1), 67–95.
13. Maxwell, M.S., Restrepo, M., Henderson, S.G., & Topaloglu, H. (2010). Approximate dynamic programming for ambulance redeployment. INFORMS Journal on Computing, 22(2), 266–281.
14. Erkut, E., Ingolfsson, A., & Erdogan, G. (2008). Ambulance location for maximum survival. Naval Research Logistics, 55(1), 42–58.
15. Dean, M.D. (2008). Improving the speed of MEDEVAC using ADP. Proceedings of the Winter Simulation Conference, 1ième−2ième. IEEE.
16. Schmid, V. (2012). Solving the dynamic ambulance relocation and dispatching problem using approximate dynamic programming. European Journal of Operational Research, 219(3), 611–621.
17. Alanis, R., Ingolfsson, A., & Kolfal, B. (2013). A Markov chain model for an EMS system with repositioning. Production and Operations Management, 22(1), 216–231.
18. Jagtenberg, C.J., Bhulai, S., & van der Mei, R.D. (2015). An efficient heuristic for real-time ambulance redeployment. Operations Research for Health Care, 4, 27–35.
19. Repede, J.F., & Bernardo, J.J. (1994). Developing and validating a decision support system for locating emergency medical vehicles in Louisville, Kentucky. European Journal of Operational Research, 75(3), 567–581.
20. Ingolfsson, A. (2013). EMS planning and management. In Operations Research and Health Care Policy (pp. 105–128). Springer.
21. Bertsimas, D., & Kallus, N. (2020). From predictive to prescriptive analytics. Management Science, 66(3), 1025–1044.
22. Chen, X., Hu, P., Shum, S., & Zhang, Y. (2020). Dynamic stochastic inventory management with reference price effects. Operations Research, 64(6), 1529–1536.
23. Sutton, R.S., & Barto, A.G. (2018). Reinforcement Learning: An Introduction (2nd ed.). MIT Press.
24. Balcik, B., & Beamon, B.M. (2008). Facility location in humanitarian relief. International Journal of Logistics, 11(2), 101–121.
25. Rawls, C.G., & Turnquist, M.A. (2010). Pre-positioning of emergency supplies for disaster response. Transportation Research Part B, 44(4), 521–534.
26. Sheu, J.B. (2007). An emergency logistics distribution approach for quick response to urgent relief demand in disasters. Transportation Research Part E, 43(6), 687–709.
27. Yi, W., & Özdamar, L. (2007). A dynamic logistics coordination model for evacuation and support in disaster response activities. European Journal of Operational Research, 179(3), 1177–1193.
28. Hssaisoune, M., Bouchaou, L., Sifeddine, A., & Bouimetarhan, I. (2020). Moroccan groundwater resources and evolution with global climate changes: Assessing groundwater vulnerability. Geosciences, 10(2), 36.
29. Boudhar, A., Boulet, H., Rhazi, L., Wafo, E., & Fakir, Y. (2021). Flood monitoring in the Gharb plain using Sentinel satellite imagery and hydrological modelling. Remote Sensing, 13(9), 1724.
30. Breiman, L. (2001). Random forests. Machine Learning, 45(1), 5–32.
31. Friedman, J.H. (2001). Greedy function approximation: A gradient boosting machine. Annals of Statistics, 29(5), 1189–1232.
32. Massey, F.J.(1951).The Kolmogorov-Smirnov test for goodness of fit. Journal of the American Statistical Association, 46(253),68–78.
33. Mohd Perdaus, A. F. bin, Kharie, M. F. bin, & Hasan, F. M. (2025). From Chaos to Calm: Best Practices in Disaster Relief Blueprint for Relief: Good Governance in Humanitarian Efforts. International Journal of Human and Health Sciences (IJHHS), 9(10), S14. https://doi.org/10.31344/ijhhs.v9i10.783
34. Nomani, M., & Parveen, R. . (2021). COVID-19 pandemic and disaster preparedness in the context of public health laws and policies. Bangladesh Journal of Medical Science, 20(5), 41–48. https://doi.org/10.3329/bjms.v20i5.55405
35. Dahou, B., Abidli, Z., Fellaq, SE., Rkhaila, A., Chahboune, M., Ziri, R., et Ahami, AOT. (2023). Impact de la pandémie de Covid-19 sur la santé psychologique des patients de la province de Kénitra au Maroc. Bangladesh Journal of Medical Science , 22 (4), 895-901. https://doi.org/10.3329/bjms.v22i4.67130