


DYNAMIC PREDICTION OF WAITING TIME INTERVALS

Application to Cancer Care Facilities


Yashi Huang, Arik Senderovich, Yaron Shaposhnik

Mind-IL: Data Science, Operations and Analytics (DSOA) | 10/26/2026



OVERVIEW


- Collaboration with the **Dana-Farber Cancer Institute (DFCI)**
 - 556 physicians and 1,063 registered nurses
 - 384,094 outpatient visits and 96,288 unique patients in 2023



Sources: [Dana-Farber Facts 2024](#)

OVERVIEW


- Collaboration with the **Dana-Farber Cancer Institute (DFCI)**
 - 556 physicians and 1,063 registered nurses
 - 384,094 outpatient visits and 96,288 unique patients in 2023
- Focus: **waiting time prediction**
 - Raised in patient surveys
 - Important operational problem that affects patient experience
 - We develop a solution and study broader questions around it



Sources: [Dana-Farber Facts 2024](#)


OVERVIEW

- Collaboration with the **Dana-Farber Cancer Institute (DFCI)**
 - 556 physicians and 1,063 registered nurses
 - 384,094 outpatient visits and 96,288 unique patients in 2023
- Focus: **waiting time prediction**
 - Raised in patient surveys
 - Important operational problem that affects patient experience
 - We develop a solution and study broader questions around it
- Agenda
 - Additional background
 - Research questions & answers
 - Takeaways




Sources: [Dana-Farber Facts 2024](#)

PATIENT FLOW AND DATA



PATIENT FLOW AND DATA



DATA RECORDED



Scheduling data (EHR)

Planned start times, durations, providers, and locations



Real-time locating (RTLS)

Badges track patients and staff: actual arrival, start, and end times



Data limitations


Noisy tracking (e.g., badges left behind); operational data only

RESEARCH QUESTIONS

THE PREDICTION TASK

Predict a patient's wait from check-in until entering the exam room

- **Interval:** lower and upper bound
- **Dynamic:** updated while the patient waits
- **Real-time:** scheduling + RTLS data
- **Requirements:** ~70% coverage, ~30-min width




RESEARCH QUESTIONS

THE PREDICTION TASK

Predict a patient's wait from check-in until entering the exam room

- **Interval:** lower and upper bound
- **Dynamic:** updated while the patient waits
- **Real-time:** scheduling + RTLS data
- **Requirements:** ~70% coverage, ~30-min width



RQ1

Which methods best balance the coverage-width tradeoff?

RQ2

What performance can be expected to be achieved in practice?

RQ3

What is the contribution of RTLS systems for waiting time prediction?

RQ4

Can the results be further improved?

RESEARCH QUESTIONS: ANSWERS

RQ1

Which methods best balance the coverage-width tradeoff?

**XGB with conformal prediction:
narrowest intervals on average, but not in
every arrival group**

RESEARCH QUESTIONS: ANSWERS

RQ1

Which methods best balance the coverage-width tradeoff?

XGB with conformal prediction:
narrowest intervals on average, but not in every arrival group

RQ2

What performance can be expected to be achieved in practice?

Offline: $\approx 70\%$ coverage with 20–30-min intervals for the main arrival groups (88% of visits). Deployed (RSF): 73–80% coverage with 35–41-min intervals

RESEARCH QUESTIONS: ANSWERS

RQ1

Which methods best balance the coverage-width tradeoff?

**XGB with conformal prediction:
narrowest intervals on average, but not in
every arrival group**

RQ2

What performance can be expected to be achieved in practice?

**Offline: $\approx 70\%$ coverage with 20–30-min
intervals for the main arrival groups (88%
of visits). Deployed (RSF): 73–80%
coverage with 35–41-min intervals**

RQ3

What is the contribution of RTLS systems for waiting time prediction?

**RTLS narrows intervals by ≈ 6 min;
recording arrival times alone captures
 ≈ 4 min of that gain**

RESEARCH QUESTIONS: ANSWERS

RQ1

Which methods best balance the coverage-width tradeoff?

XGB with conformal prediction: narrowest intervals on average, but not in every arrival group

RQ2

What performance can be expected to be achieved in practice?

Offline: $\approx 70\%$ coverage with 20–30-min intervals for the main arrival groups (88% of visits). Deployed (RSF): 73–80% coverage with 35–41-min intervals

RQ3

What is the contribution of RTLS systems for waiting time prediction?

RTLS narrows intervals by ≈ 6 min; recording arrival times alone captures ≈ 4 min of that gain

RQ4


Can the results be further improved?

Only marginally: an idealized optimal predictor does about as well, so much of the error is irreducible (unless more data is collected)

NESTEDSIMPY: HOW GOOD CAN ML PREDICTIONS GET?

Work in progress with **Hongyu Wang** · nestedsimpy.github.io

- **NestedSimPy**: open-source extension of SimPy (Python); at *trigger events* the simulation pauses and branches into many inner simulations from the same state
- **Unbiased branching**: remaining durations are resampled given the time already elapsed




Inner simulations launched at one trigger event (NestedSimPy output, M/M/1 queue)

NESTEDSIMPY: HOW GOOD CAN ML PREDICTIONS GET?

Work in progress with **Hongyu Wang** · nestedsimpy.github.io

- **NestedSimPy**: open-source extension of SimPy (Python); at trigger events the simulation pauses and branches into many inner simulations from the same state
- **Unbiased branching**: remaining durations are resampled given the time already elapsed



Inner simulations launched at one trigger event (NestedSimPy output, M/M/1 queue)

USE CASES



Benchmarks for ML

Measure the gap to the optimal predictor, not only rankings



Teaching analytics

Design prediction exercises in which the best achievable performance is known



Lookahead policies

Simulate alternative actions before deciding (rollout policies)

DEPLOYMENT

Select a Clinic
DF YAWKEY 7

Model v1.0 | Run #1591
Refreshed 6 minutes ago

Patient Name ↓	DFCI MRN	Appointment Date	Expected Wait Time	Upper Bound	Lower Bound	Provider Name	Department	Visit Type
		08/11/2025 15:30 PM	50 minutes	74 minutes	5 minutes		DF LYMPHOMA	ESTABLISHED PATIENT
		08/11/2025 15:00 PM	45 minutes	57 minutes	13 minutes		DF LYMPHOMA	TEACHING / STUDY
		08/11/2025 15:00 PM	26 minutes	34 minutes	6 minutes		DF MULTIPLE MYELOMA	ESTABLISHED PATIENT ON TX
		08/11/2025 14:30 PM	27 minutes	37 minutes	4 minutes		DF MULTIPLE MYELOMA	ESTABLISHED PATIENT ON TX
		08/11/2025 14:30 PM	42 minutes	50 minutes	4 minutes		DF MULTIPLE MYELOMA	ESTABLISHED PATIENT
		08/11/2025 15:30 PM	61 minutes	75 minutes	26 minutes		DF MULTIPLE MYELOMA	ESTABLISHED PATIENT ON TX
		08/11/2025 11:00 AM	170 minutes	311 minutes	30 minutes		DF LYMPHOMA	CONSULT
		08/11/2025 14:30 PM	18 minutes	25 minutes	3 minutes		DF MULTIPLE MYELOMA	ESTABLISHED PATIENT
		08/11/2025 14:30 PM	10 minutes	16 minutes	2 minutes		DF LYMPHOMA	ESTABLISHED PATIENT

Select a Clinician

Patient Name ↓	DFCI MRN	Appointment Date	Expected Wait Time	Upper Bound	Lower Bound
		08/11/2025 16:00 PM	46 minutes	65 minutes	8 minutes
		08/11/2025 15:00 PM	29 minutes	33 minutes	3 minutes

Web application used by clinical assistants (identifiers hidden)

- First version based on RSF
- Real-time pipeline from EHR and RTLS data
- Model size cut from 20 GB to 2 GB
- Display refreshes every 5 minutes
- Data-drift and performance monitoring
- Half-intervals when too wide: “at least X min” (or “uncertain”)

Jun–Aug 2025

On-site development

Development and testing on site; by the end, used for “escalations”

Feb 27, 2026

Phase 1: “quiet” pilot

A month of staff-only access: is the app stable, used with every patient, easy to interpret, and useful?

Apr 17, 2026

Phase 2: sharing with patients

Clinic Assistants share each patient’s predicted wait after vitals are taken

Aug 31, 2026

Phase 3: waiting-room display

Predictions shown on a TV application in the waiting room

TAKEAWAYS

RQ1

Which methods best balance the coverage-width tradeoff?

XGB with conformal prediction: narrowest intervals on average, but not in every arrival group

RQ2

What performance can be expected to be achieved in practice?

Offline: $\approx 70\%$ coverage with 20–30-min intervals for the main arrival groups (88% of visits). Deployed (RSF): 73–80% coverage with 35–41-min intervals

RQ3

What is the contribution of RTLS systems for waiting time prediction?

RTLS narrows intervals by ≈ 6 min; recording arrival times alone captures ≈ 4 min of that gain

RQ4

Can the results be further improved?

Only marginally: an idealized optimal predictor does about as well, so much of the error is irreducible (unless more data is collected)

PRACTICAL TAKEAWAYS



Building the model

- **Use the schedule:** slot order, gaps, provider lateness
- **Tune the quantiles:** asymmetric bounds cut ≈ 5 min of width
- **Calibrate per arrival group:** a small lookup table keeps coverage on target
- **Simple can work:** XGB is narrowest on average, but check late arrivals



Communicating predictions

- **Intervals, not points:** much of the uncertainty is irreducible
- **Half-intervals when too wide:** “at least X min” or “uncertain”
- **Keep updating:** coverage holds while patients wait; widths don't shrink
- **Imperfect estimates still help:** 10 of 13 patients found them helpful



Deploying it

- **Start without tracking:** arrival times give $\approx 2/3$ of the RTLS gain
- **Live $\neq$ back-test:** monitor coverage and width in production
- **Engineering matters:** model size, refresh, drift monitoring, versioning
- **XAI helped developers, not CAs:** it validated the model; CAs (arguably) lack the time and expertise to analyze predictions
- **Requirements evolve:** developing a single general model is beneficial
- **Surprisingly expensive:** staff hours and cloud computing

A photograph of a university campus. In the foreground, a large, mature tree with green leaves dominates the right side. Its shadow is cast on the grass. In the middle ground, the word "MELIORA" is spelled out in large, white, three-dimensional block letters on a green lawn. In the background, there are university buildings and more trees. A semi-transparent blue box is overlaid on the right side of the image.

THANK YOU

Questions?