
November 15, 2022
With a Little Help From My Friends
Tools and insights for developing and deploying algorithms in the hospital

November 15, 2022

With a Little Help From My Friends

How we got our model out of one person's computer and deployed it to the hospital

Agenda

- A bit about DSAA
- Where we started vs where we are now
- What do you need for **development**?
- What do you need for **deployment**?

Introduction

A little bit about DSAA

- Hello! I'm **Chloe** and I'm a data scientist for the DSAA team at Unity Health Toronto.

- **Data Science and Advanced Analytics (DSAA)** is a healthcare data analytics group at **Unity Health Toronto (UHT)**.

UNITY HEALTH TORONTO

Data Science and Advanced Analytics

A little bit about DSAA

- Suit the needs of the hospital, our collaborators, and our partners to *make better decisions*, *increase hospital efficiency*, and *improve patient care and patient outcomes*.
- DSAA works with *clinicians* and *administrative decision-makers* to develop and deploy solutions.

A little bit about DSAA

- Ranges of solutions: statistics, artificial intelligence, machine learning, and optimization (e.g., operations research).
- Currently more than 40 active solutions at Unity Health:
  - Predicting patient outcomes for enhanced clinical management
  - Planning for hospital bed capacity
  - Medical imaging AI tools
  - Assignment/scheduling

A little bit about DSAA

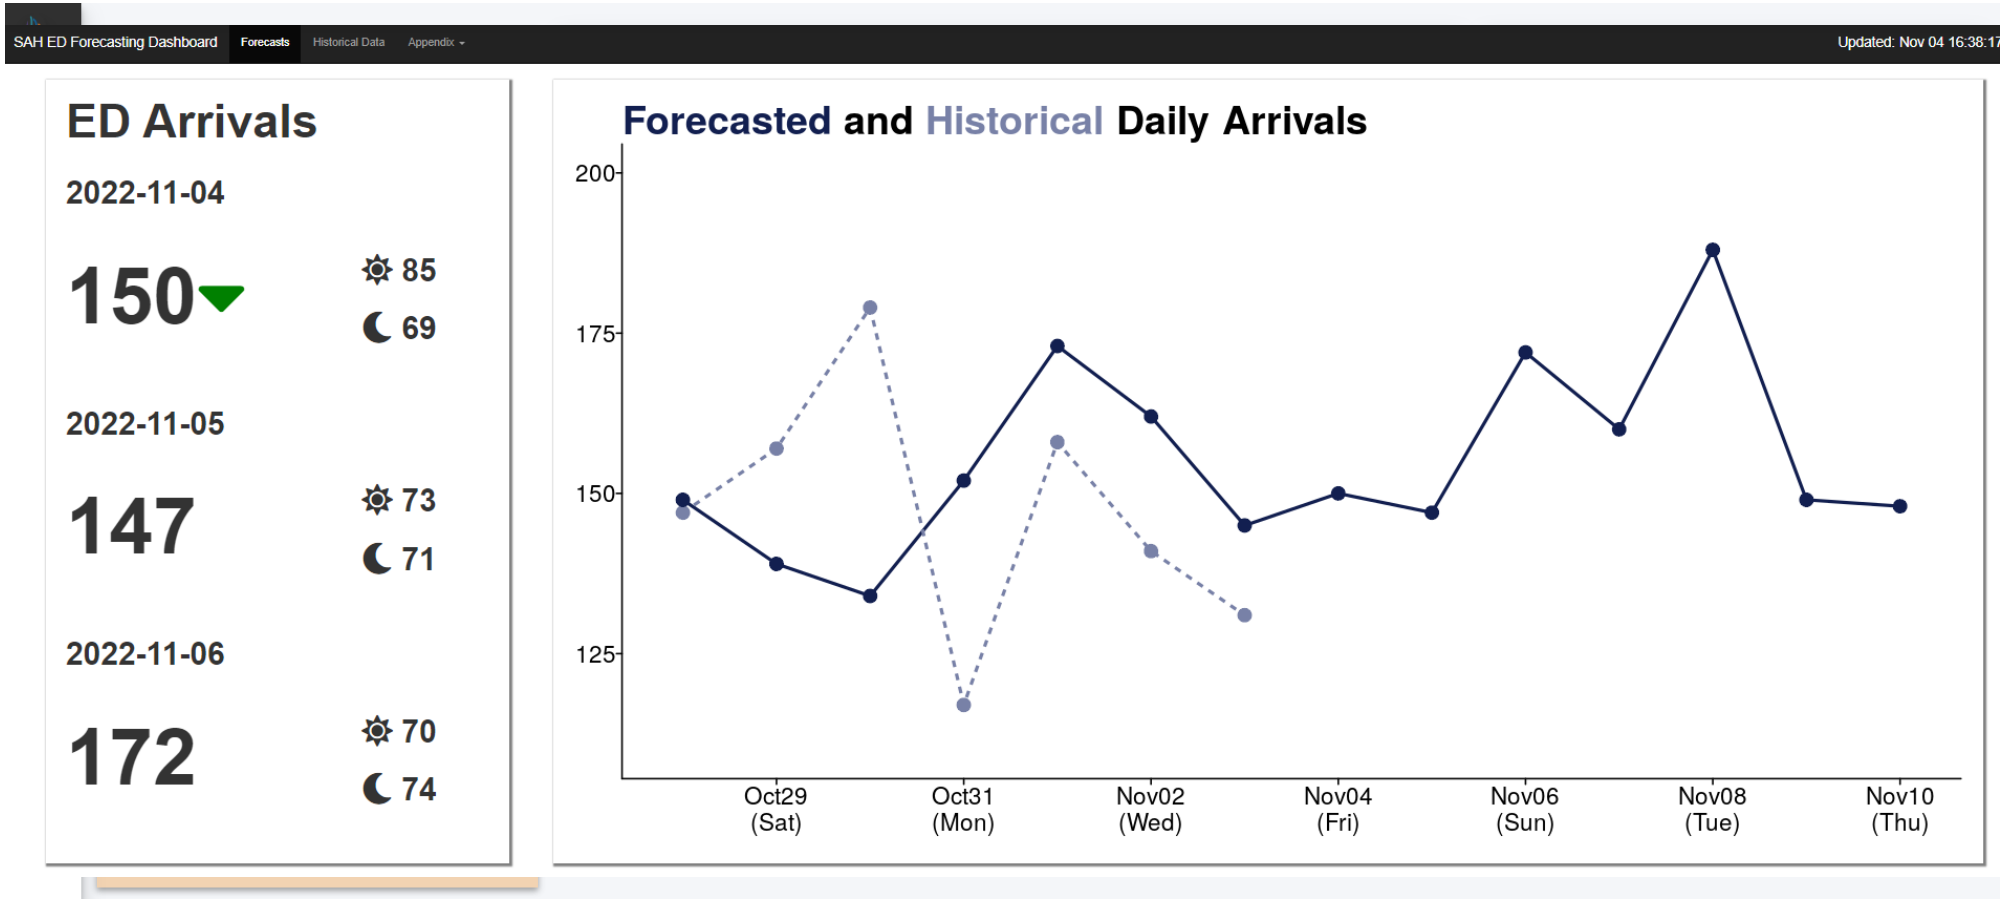

A little bit about DSAA

A little bit about DSAA

- DSAA is split into 4 teams:
  - Data Integration and Governance team
  - Advanced Analytics team
  - Product Development team
  - Support team

A successful deployment!

- In August 2020, we deployed **CHARTwatch**, an early warning system for detecting patient deterioration.
- CHARTwatch was deployed to the General Internal Medicine (GIM) ward.
- *Aside: Why is it called "CHARTwatch"?*

CHARTwatch

- We developed a model to predict which patients are at risk of deterioration:
  - Transfer to the Intensive Care Unit (ICU)
  - Death
  - Transfer to the Palliative Care unit
- The model was trained on ~20,000 patient visits consisting of: laboratory values, vital measurements, and demographics.
- The model output is a risk group: **High** vs **Medium** vs **Low** risk group

CHARTwatch

Model predictions were delivered to different end-users.

- Email to charge nurses
- Email to the Palliative Care team

To Palliative team: The following GIM patients were identified by CHARTwatch as being at high risk of dying or requiring ICU in the next 48 hours.

	ENCOUNTER_NUM	LOCATION	BED	MRN	PATIENT_NAME	AGE	SEX	TEAM	PHYSICIAN	STATUS	ISOLATION_STATUS	ALARM_TS
1		14CG	K040-C2					Team Medicine C		High	Exposed COVID-19	22:01:17
2		8CS	844C-2	2				Team Medicine B		High	Contact Precautions	07:40:54
3		14CU	K010-04					Team Medicine A		Medium		22:01:17
4		14CG	K718-G1	2				Team Medicine E		Medium	Confirmed COVID-19	15:01:08
5		14CG	K728-G1	3				Team Medicine D		Medium	Suspected COVID-19	15:01:08

CHARTwatch

Model predictions were delivered to different end-users.

- Updates to front-end tool

St. Michael's
Inspired Care. Inspiring Science.

eDischarge Summary | eSign-out | GIM eAdmit

Team Help Welcome [redacted] | Logout

Teams»General Internal Medicine Team E

Sign-out List | Team Profile | View/Print Full | View/Print On-Call | Signed Off

Sign-out List

Add New Patient to Sign-out List

Patient MRN Search

Last Name	First Name	MRN	Gender	DOB	Age	Encounter #	Nursing Unit-Room-Bed	Code Status	CHART WATCH	Admission Date	House Staff	Detail
TEST-MOTHER	BABY-BOY	[redacted]	M	Sep 30, 2014	4y	[redacted]	15NB - L026 - 2N			Sep 30, 2014 08:00:00		Sign-out Detail
CPOE	Test2	[redacted]	F	Jan 01, 1960	59y	[redacted]	3B - - 368B1		HIGH	Jun 03, 2018 14:00:00		Sign-out Detail
[redacted]							3B - 374B - 2		low	Feb 16, 2018 15:40:00		Sign-out Detail
[redacted]							7CC - 708C - 02			Oct 22, 2018 11:12:00		Sign-out Detail
[redacted]							7CC - 716C - 1	Full Code	medium	Oct 22, 2018 11:11:00	8	Sign-out Detail
[redacted]							7CCV - 744C - 1			Jul 03, 2015 11:24:00		Sign-out Detail
[redacted]							7CEL - 704C - 2			Mar 26, 2018 15:33:00		Sign-out Detail
[redacted]							7CEL - 708C - 2	No CPR: Advanced Life Support		Mar 26, 2018 15:33:00		Sign-out Detail

Add New Patient to Sign-out List

CHARTwatch

Model predictions were delivered to different end-users.

- Alerts sent to phones

Read more about it:

- “Preparing a Clinical Support Model for Silent Mode in General Internal Medicine”
- “Implementing Machine Learning in Medicine”
- “From Compute to Care: Lessons Learned from Deploying an Early Warning System into Clinical Practice”

Back to the future

Back to the future

Before

- Scripts running from one person's laptop
- No logging
- No development environments

Now

- Service accounts!
- Log all the things!
- Staging vs production environments!
- …. and more!

How did we get from “Before” to “Now”?

- Tools for **development**
- Tools for **deployment**

What do you need for *development*?

Connections to databases

- Various data systems in the hospital.
- We developed `chartdb`, an internal R package to interact with hospital databases.
- All connection functions follow the same pattern.

Connections to databases

```
con_a <- chartdb::connect_databaseA(username = ..., password = ...)
con_b <- chartdb::connect_databaseB(username = ..., password = ...)
con_edw <- chartdb::connect_edw(username = ..., password = ...)
con_soarian <- chartdb::connect_soarian(username = ..., password = ...)
con_mak <- chartdb::connect_mak(username = ..., password = ...)
con_syngo <- chartdb::connect_syngo(username = ..., password = ...)
```

```
while(try < retries){
    con <- CHART_connect_odbc(source_db = "EDW",
                              username = username,
                              password = password)
    if(class(con) == "NetezzaSQL"){
      try <- retries + 1
    } else if (!"NetezzaSQL" %in% class(con) & try < retries){
      warning("EDW connection failed. Retrying...")
      try <- try + 1
      Sys.sleep(retry_wait)
    } else {
      try <- try + 1
      warning("EDW connection failed")
    }
  }
```

Connections to databases

- `chartdb` also offers other utility functions:

```
con_edw <- chartdb::connect_edw(username = ..., password = ...)
adt <- get_patient_adt(con_edw)
```

- ADT = Admit / Discharge / Transfer

An environment that lets multiple people collaborate on a project

An environment that lets multiple people collaborate on a project

- We started using renv.

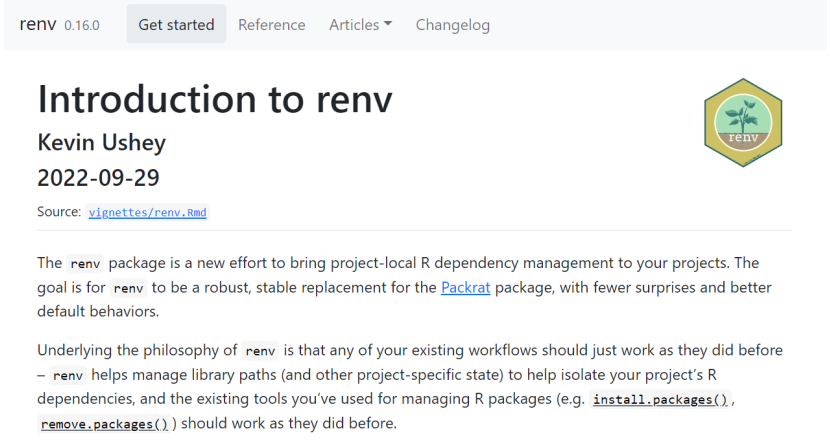

- renv is an R package for R dependency management.

An environment that lets multiple people collaborate on a project

```
very_awesome_project/
|- .github/
|- R/
|- renv/
|- tests/
|- DESCRIPTION
|- NAMESPACE
|- README.md
|- .gitignore
|- renv.lock
```

An environment that lets multiple people collaborate on a project

```
  ...
  "Packages": {
    ...
    "dplyr": {
      "Package": "dplyr",
      "Version": "1.0.7",
      "Source": "Repository",
      "Repository": "RSPM"
    },
    "lubridate": {
      "Package": "lubridate",
      "Version": "1.7.10",
      "Source": "Repository",
      "Repository": "RSPM"
    },
    ...
  }
```

What do you need for *development*?

✓ Connections to databases

✓ An environment that lets multiple people collaborate on a project

Functions and utilities that you re-use

- … package-based development!

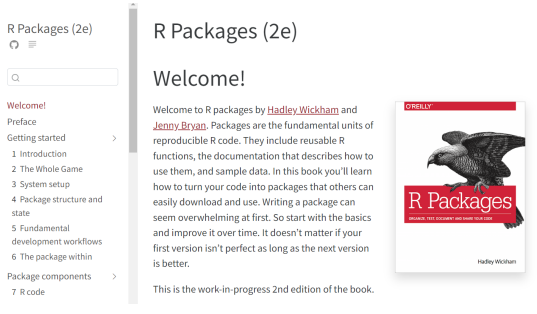

- Why write a package?
  - Share code/knowledge with others
  - No more copy-pasting

Functions and utilities that you re-use

```
very_awesome_project/
|- .github/
|- R/
|- renv/
|- tests/
|- DESCRIPTION
|- NAMESPACE
|- README.md
|- .gitignore
|- renv.lock
```

Functions and utilities that you re-use

```
Package: chartwatch
Type: Package
Title: Utilities For CHARTwatch Project
Version: 1.11.10
Author: DSAA
Maintainer: Chloe Pou-Prom <Chloe.Pou-Prom@unityhealth.to>
Description: This package contains utility functions for CHARTwatch,
    the General Internal Medicine Early Warning System.
License: MIT + file LICENSE
Encoding: UTF-8
LazyData: true
Suggests:
    testthat
RoxygenNote: 7.1.1
Imports:
  lubridate,
  ...
...
```

What do you need for *development*?

✓ Connections to databases

✓ An environment that lets multiple people collaborate on a project

✓ Functions and utilities that you re-use

Environments! Environments! Environments!

- We work with `development`, `staging` and `production` environments.
- The `development` environment:
  - Local computer
  - The `development` server: GPUs, works with `chartdb`

Environments! Environments! Environments!

- The `staging` environment is as close to the “real” deployment environment as possible.
  - When we need to make updates to CHARTwatch, we first deploy it to the `staging` environment.
- The `production` environment is where things actually get deployed.

What do you need for *development*?

✓ Connections to databases

✓ An environment that lets multiple people collaborate on a project

✓ Functions and utilities that you re-use

✓ An environment that is similar to the production environment

What do you need for *deployment*?

Pause... story time

PAUSE Reflect on CHARTwatch's deployment

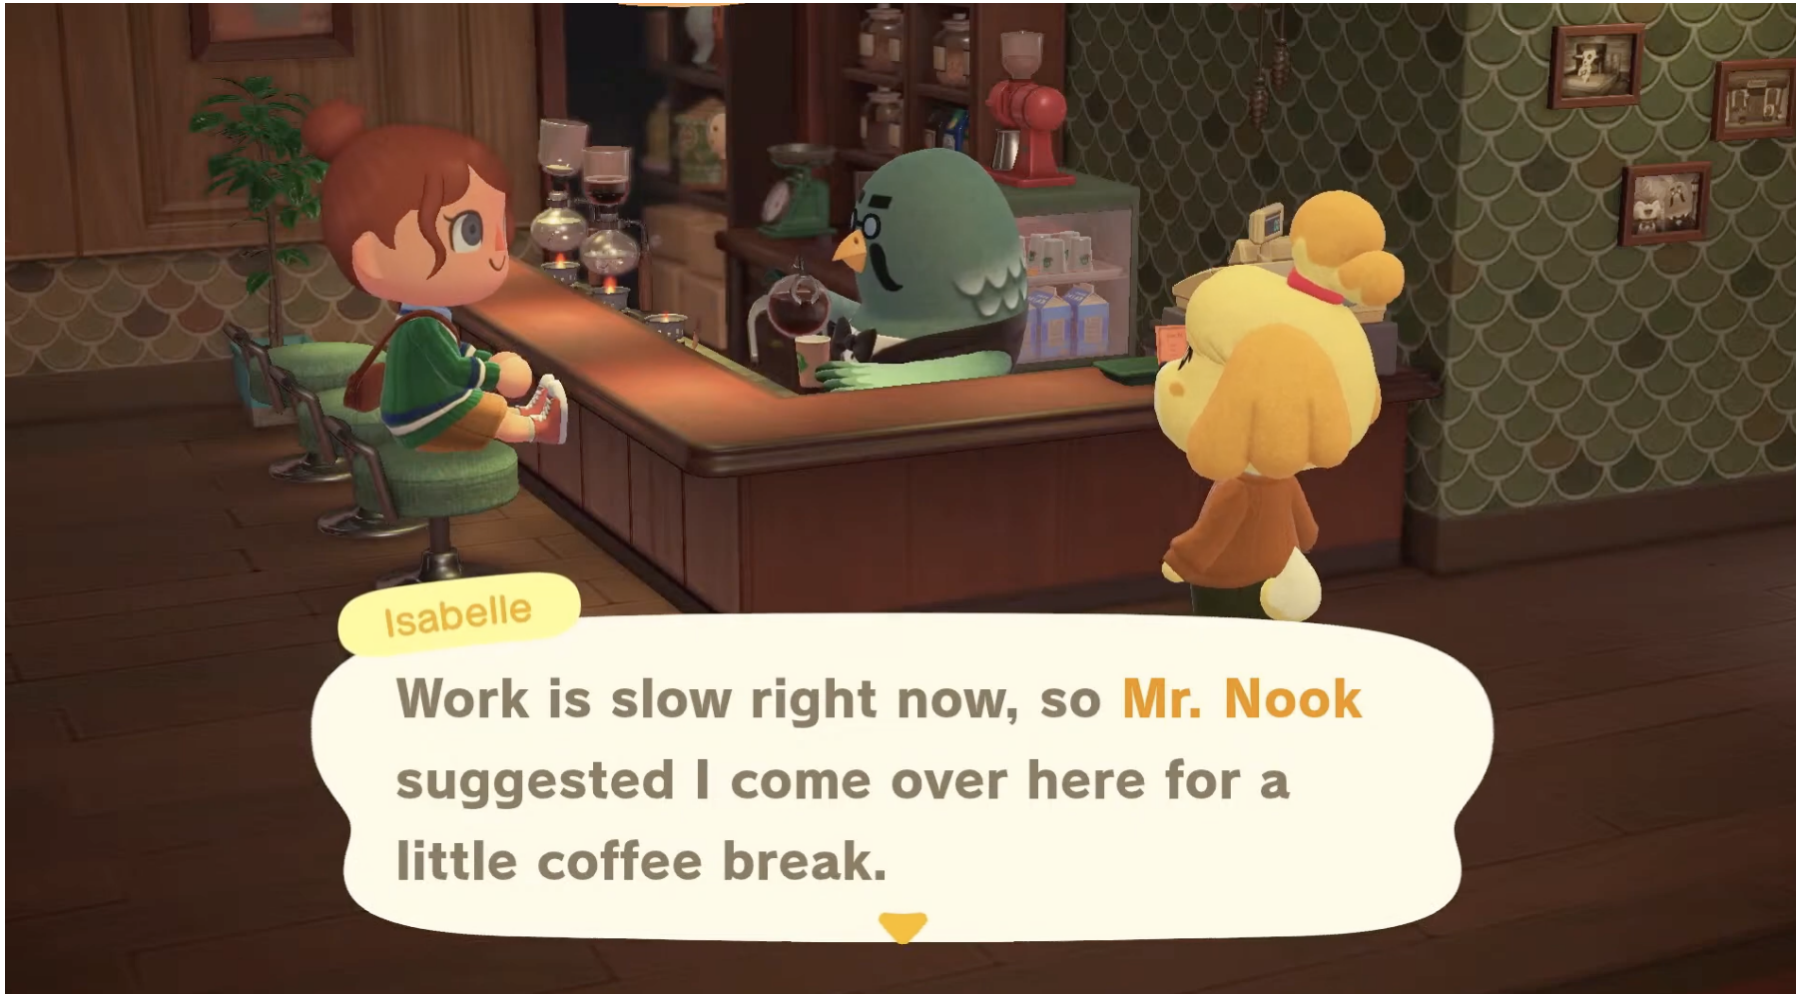

Pause... story time

- CHARTwatch was deployed on August 2020.
- It was first silently deployed in November 2019.
- Originally, CHARTwatch was supposed to go live in early 2020, but the pandemic affected our plans…

Pause... story time

00-data-extraction	Add Order Status History (#57)	3 years ago
00-setup	CHARTwatch v0.4 (#72)	3 years ago
01-data-preprocessing	CHARTwatch v0.4 (#72)	3 years ago
02-nurse-notes-processing	README updates (#74)	3 years ago
03-ensembling	CHARTwatch v0.4 (#72)	3 years ago
04-alarming	CHARTwatch v0.4 (#72)	3 years ago
05-email	Show isolation status in email (#76)	3 years ago
06-cleanup	CHARTwatch v0.4 (#72)	3 years ago
.gitignore	Update gitignore	3 years ago
README.md	Update vs-research mount instructions	3 years ago
error_check.py	Send email on failure (#30)	3 years ago
gim_ews_production.Rproj	Email stuff	3 years ago

Pause... story time

This is what the first attempt at deployment looked like for CHARTwatch:

- A mixture of Python and R scripts
- A CRON job…….. that calls different bash scripts…….. that calls different Python/R scripts

```
55 7 * * TUE sh /home/smhuser/gim-ews/code/run_generate_predictions.sh
0 10 * * TUE sh /home/smhuser/gim-ews/code/run_email_palliative.sh
20 15 * * TUE sh /home/smhuser/gim-ews/code/run_full_pipeline.sh
50 22 * * TUE sh /home/smhuser/gim-ews/code/run_generate_predictions.sh
0 3 * * WED sh /home/smhuser/gim-ews/code/run_email_charge_nurses.sh
```

What do you need for *deployment*?

Authentication

- Our deployed applications run on **RStudio Connect**

Authentication

- Our deployed applications run on ~~**RStudio Connect**~~ **Posit Connect**.

Authentication

- Our deployed applications run on ~~**RStudio Connect**~~ **Posit Connect**.
- ~~RStudio Connect~~ Posit Connect connects to the hospital's Active Directory.

Authentication

- What this means:
  - Users can authenticate using their hospital username and password!
  - For developers: don’t need to keep track of an extra server username/password
  - For end-users: to access application, they log in with their hospital credentials
  - We can use existing Active Directory groups to manage permissions

Scheduling scripts

- ~~RStudio Connect~~ Posit Connect allows us to schedule scripts.
- We run our applications from service accounts.
- Automatic scheduling + service account = deployments don’t rely on one person

Scheduling scripts

```
55 7 * * TUE sh /home/smhuser/gim-ews/code/run_generate_predictions.sh
0 10 * * TUE sh /home/smhuser/gim-ews/code/run_email_palliative.sh
20 15 * * TUE sh /home/smhuser/gim-ews/code/run_full_pipeline.sh
50 22 * * TUE sh /home/smhuser/gim-ews/code/run_generate_predictions.sh
0 3 * * WED sh /home/smhuser/gim-ews/code/run_email_charge_nurses.sh
```

What do you need for *deployment*?

✓ Authentication

✓ Scheduling

To Know when there's a downtime

- Alerts!

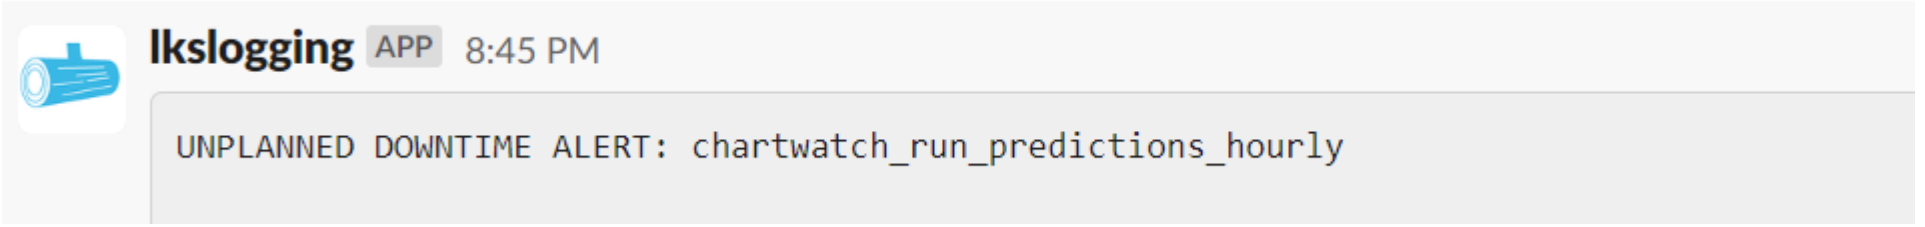

- We created `jarvis`, an R package for helping us monitor our production applications.

To Know when there's a downtime

- We use `jarvis` to send email alerts.

```
1  jarvis::send_email(
2    type = "ERROR",
3    project = "COVID Dashboard"
4  )
```

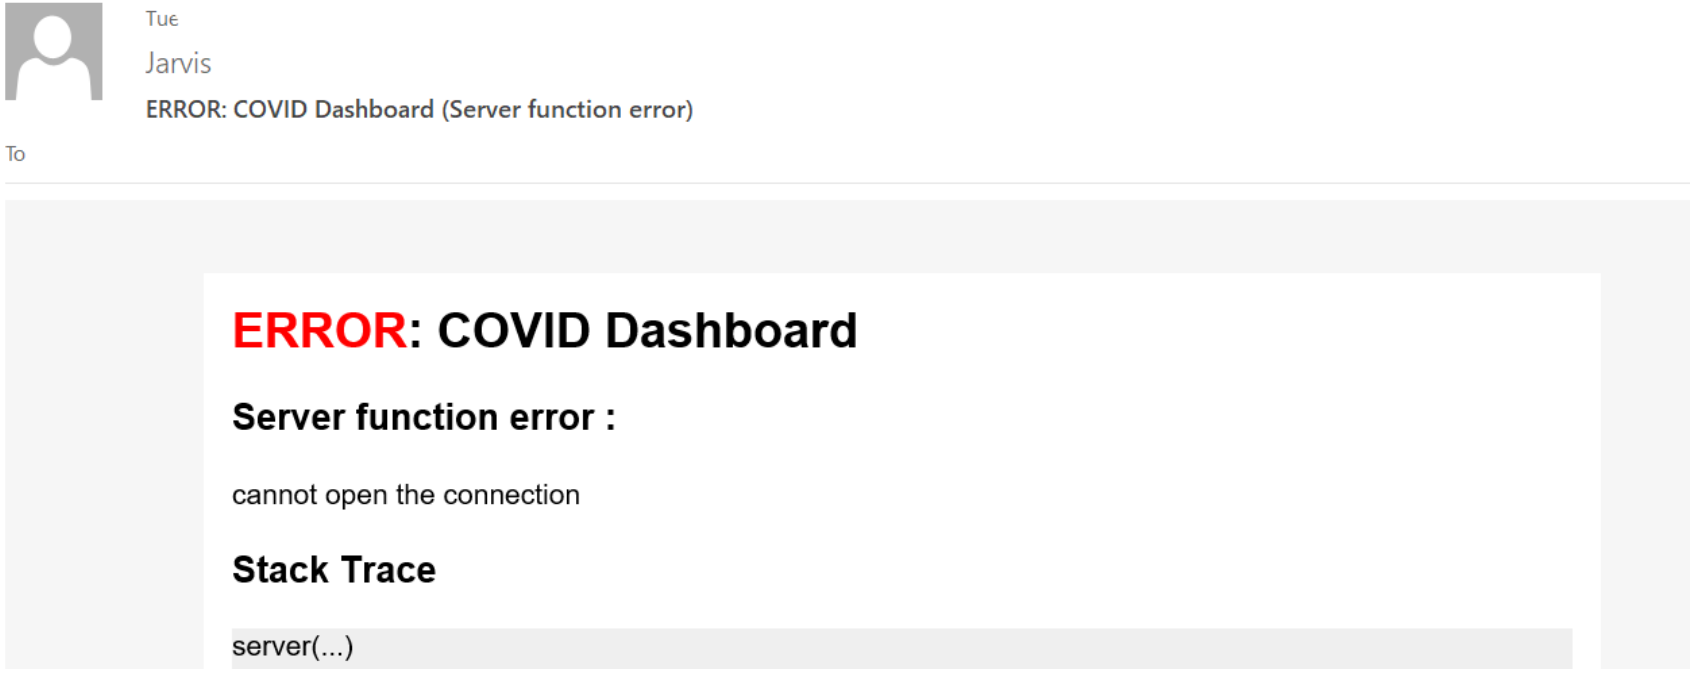

To Know when there's a downtime

- We use `jarvis` to send Slack messages.

```
jarvis::send_slack(
    channel = "#general",
    message = "Hello World!",
    level = "INFO",
    slack_api_key = Sys.getenv("SLACK_API_KEY")
)
```

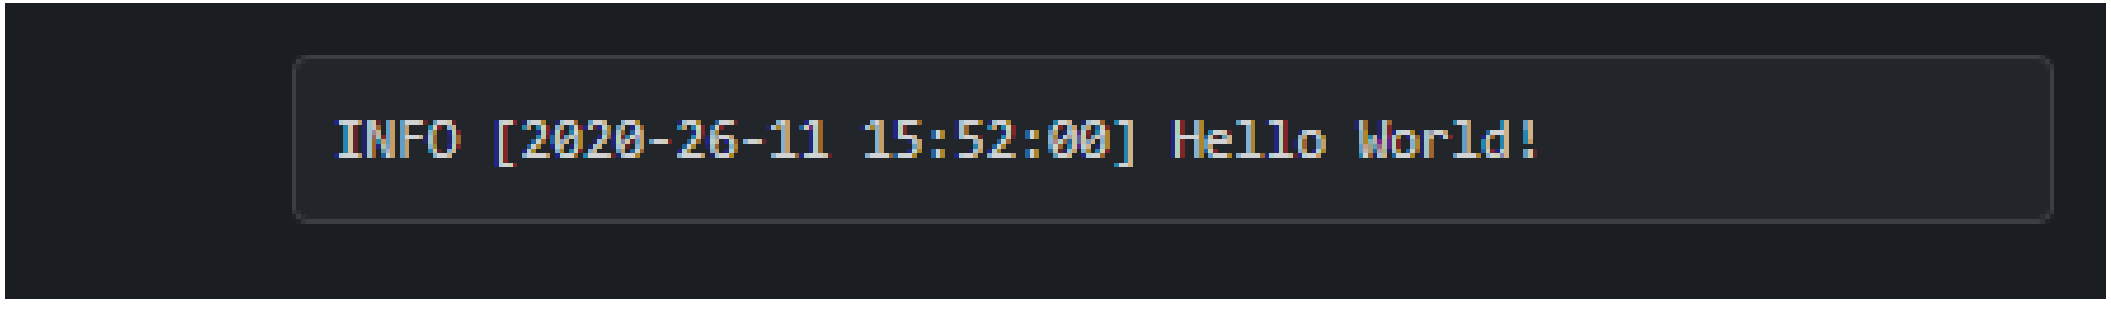

To Know when there's a downtime

- We use `jarvis` to check the health of systems we depend on.

```
# Checking a suite of dependencies (eg. EDW, Soarian, and vs-research)
jarvis::check_health(c("edw", "soarian", "vs-research"))
#  edw     soarian vs-research
# TRUE        TRUE        TRUE

# Checking a specific database
jarvis::check_db("edw")
#  edw
# TRUE

# Checking a specific mounted filesystem
jarvis::check_mount("vs-research")
#  vs-research
# TRUE
```

To Know when there's a downtime

- We have downtime protocols.

SMH Unplanned Downtime – CHARTwatch	
When?	
Why?	SMH network issues.
Who is impacted?	All users of: • CHARTWatch
What do you need to do?	Resume clinical operations. CHARTWatch alerts will not be sent until services are restored.
For any other questions or concerns regarding this issue, or if you feel you are experiencing problems, please contact the Helpdesk at: ***PHC:*** ***SJHC:*** ***SMH:*** *Please do not respond directly to this email, as the originating e-mail account is not monitored.*	
Notification Colour Codes	
Planned Downtime	**Unplanned Downtime** / **Resolved Downtime** / **Service Interruption**

What do you need for *deployment*?

✓ Authentication

✓ Scheduling

✓ To know when there’s a downtime

A secure way to download internally-developed packages

A secure way to download internally-developed packages

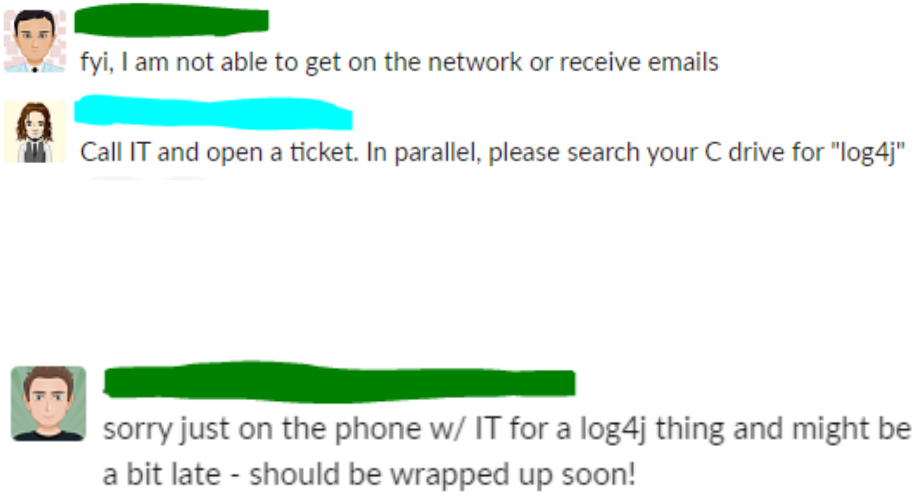

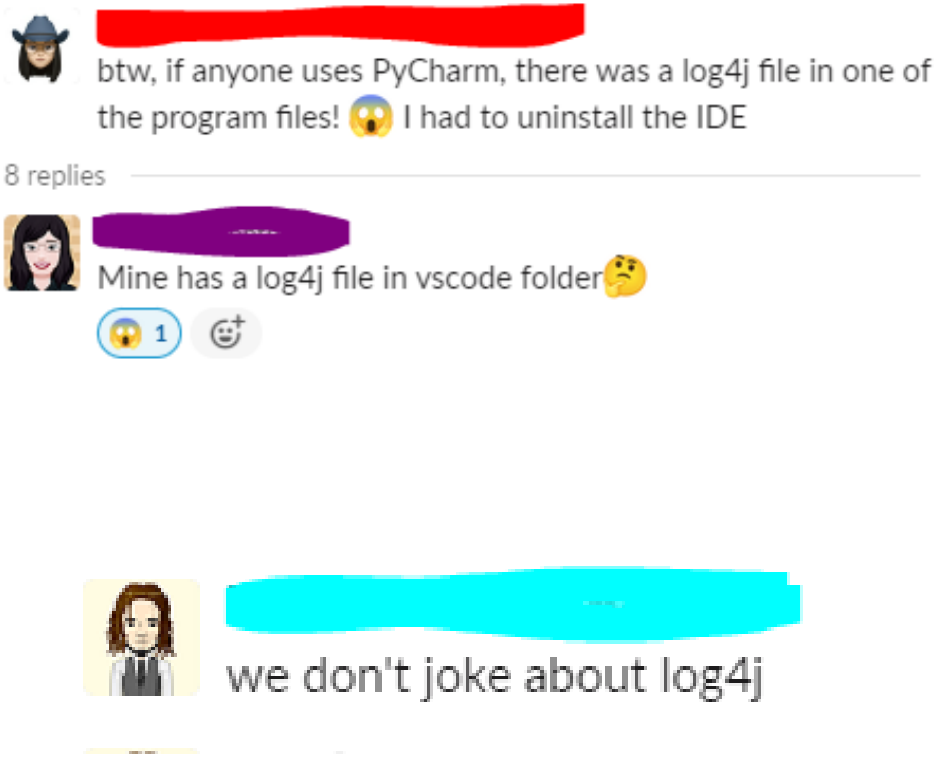

A secure way to download internally-developed packages

- `log4j` is a Java-based logging utility.
- In December 2021, it was discovered that the "`log4j` flaw" could allow malicious users to access internal networks.

A secure way to download internally-developed packages

Log4j is a pervasive vulnerability. Update your devices now

A hole in a popular piece of code is an open window for criminals.

Log4j software flaw 'endemic,' new cyber safety panel says

A secure way to download internally-developed packages

- We need to limit who can access the hospital network.
- ~~RStudio Package Manager (RSPM)~~ Posit Package Manager is a repository management server.
- We can download packages while being disconnected from the Internet.

A secure way to download internally-developed packages

- Remember when we mentioned `renv`?

```
...
"Packages": {
  ...
  "dplyr": {
    "Package": "dplyr",
    "Version": "1.0.7",
    "Source": "Repository",
    "Repository": "RSPM"
  },
  "lubridate": {
    "Package": "lubridate",
    "Version": "1.7.10",
    "Source": "Repository",
    "Repository": "RSPM"
  },
```

What do you need for *deployment*?

✓ Authentication

✓ Scheduling

✓ To know when there's a downtime

✓ A secure way to download internally-developed packages

An implementation plan


2
A
B
C
D
E
F
1
2
3
4
5
6
7

An implementation plan

- CHARTwatch predictions are delivered to different end-users:
  - Email to charge nurses
  - Email to the Palliative Care team
  - Updates to front-end tool
  - Alerts sent to phones

An implementation plan

- CHARTwatch predictions are delivered to different end-users:
  - Email to **charge nurses**
  - Email to the **Palliative Care team**
  - Updates to front-end tool (**IT**)
  - Alerts sent to phones (**IT**, **residents**, **physicians**)

An implementation plan

- CHARTwatch was deployed by an implementation team consisting of many people.
  - General Internal Medicine (GIM)
  - Intensive Care Unit (ICU)
  - Palliative Care
  - Clinical Informatics
  - IT
  - Data science team

An implementation plan

What happens after CHARTwatch flags a patient as being High risk?

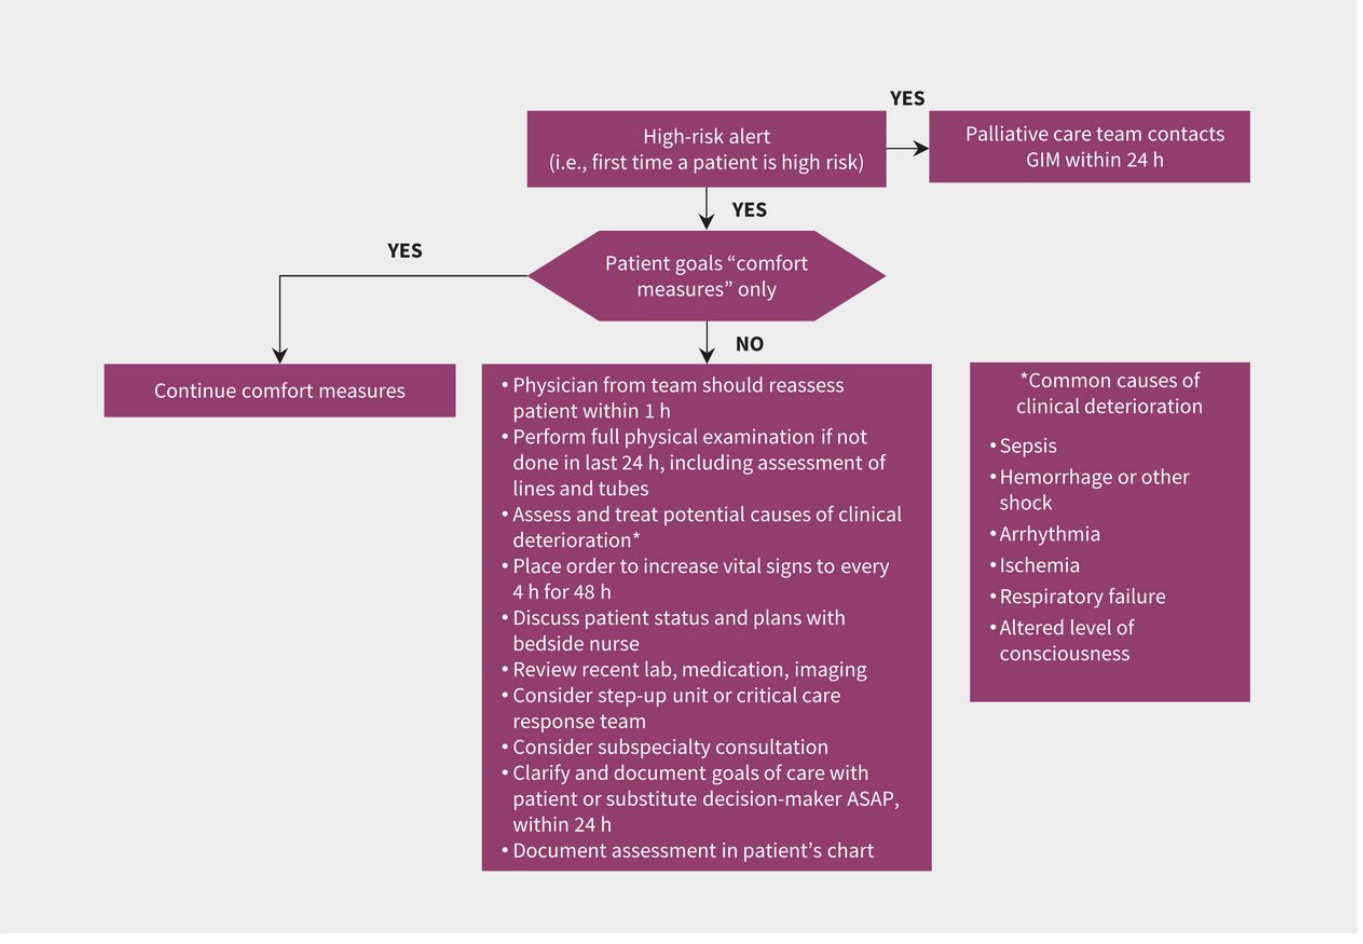

An implementation plan

What happens after CHARTwatch flags a patient as being High risk?

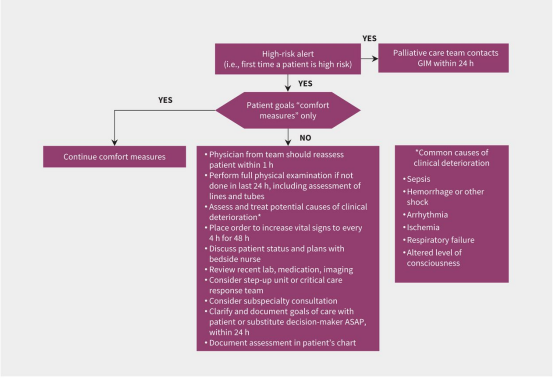

- Consider existing resources
- Alerting notifications fit within existing processes
  - Timing of emails to charge nurses
- Clinical pathway
  - Time targets
  - Leave room for clinical judgment

An implementation plan

- Silent deployment period

- Change to how **troponin** is measured.
- Changes due to deploying at the beginning of the pandemic.

An implementation plan

- Silent deployment period

- Catching bugs!
  - How are missing values represented in R?
    - NA (not available)
  - What's the chemical element for sodium?
    - NA = sodium

An implementation plan

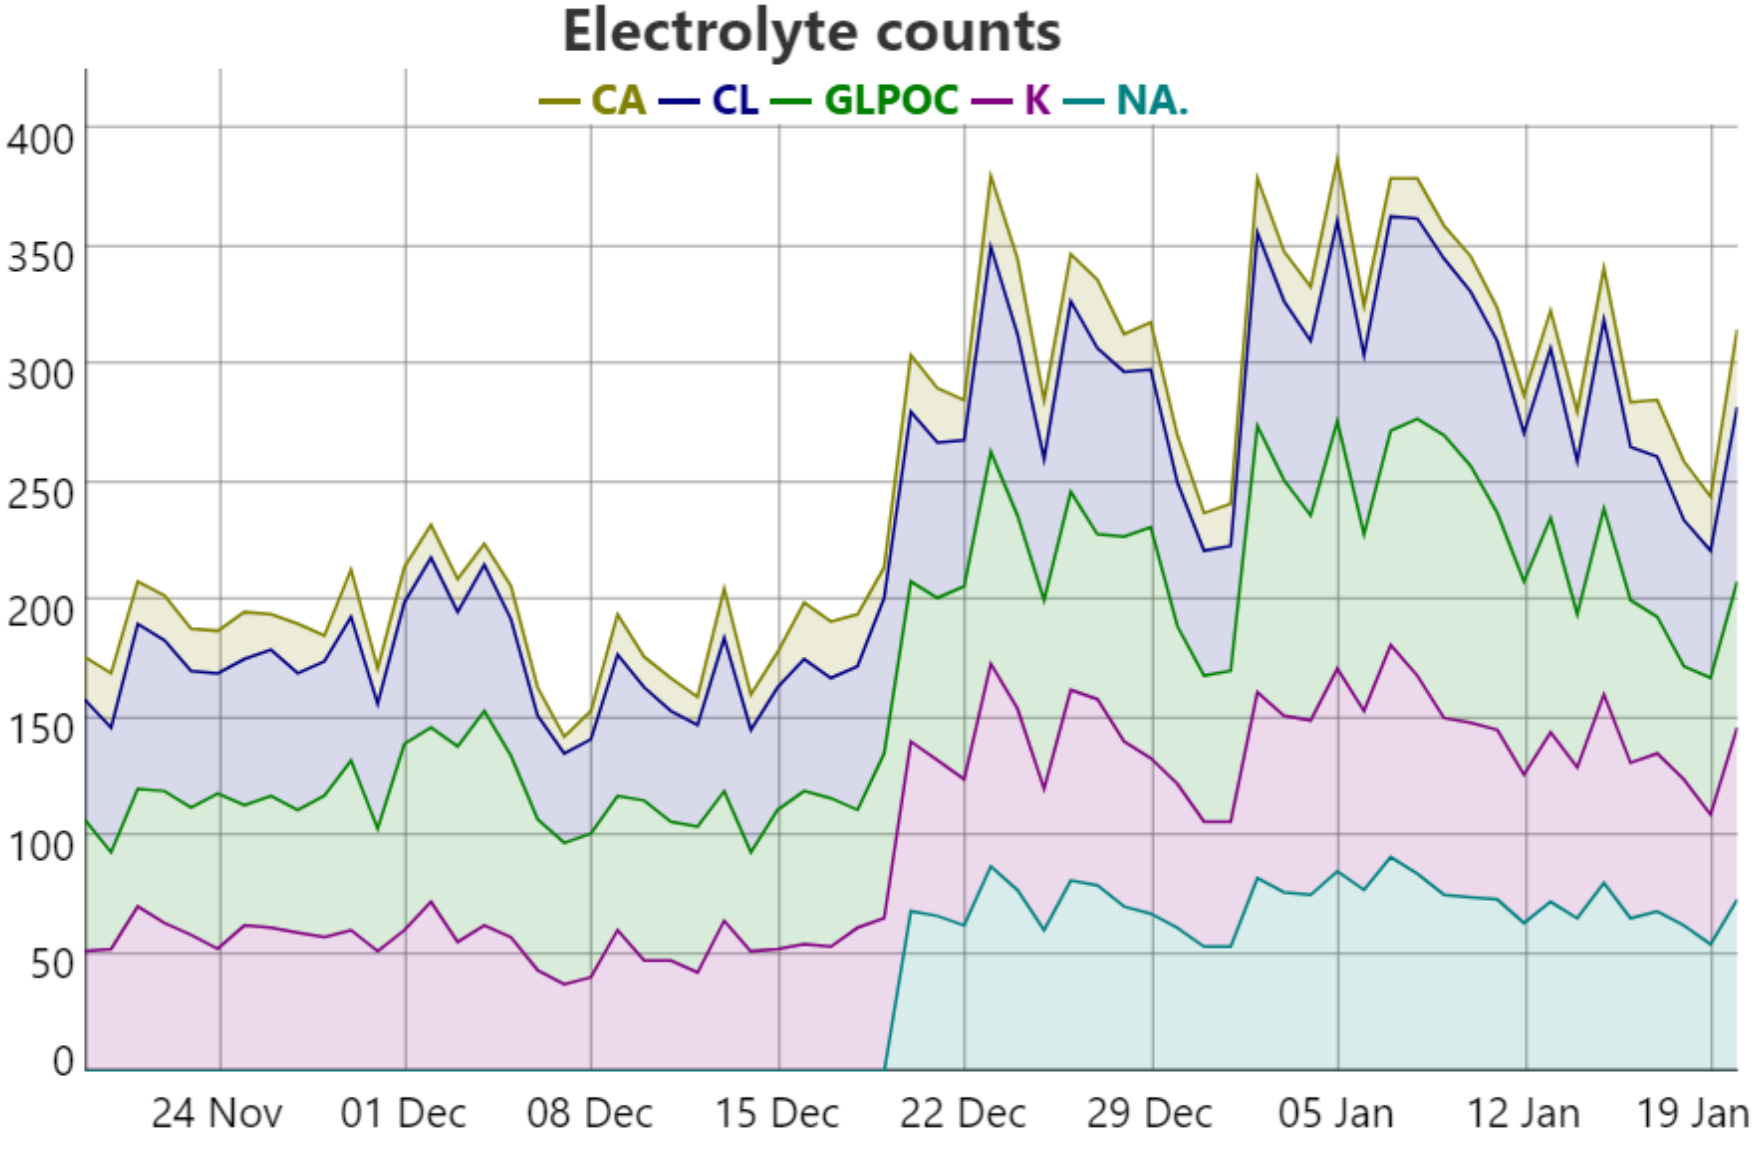

An implementation plan

- **Pilot phase**
  - CHARTwatch was first deployed to two GIM teams.
  - Weekly meetings
- End-user engagement is important!
  - Developing the clinical pathway
  - Developing education and training processes

What do you need for *deployment*?

✓ Authentication

✓ To know when there's a downtime

✓ A secure way to download internally-developed packages

✓ An implementation plan

Something about monitoring...?

Something about monitoring...?

Something about monitoring...?

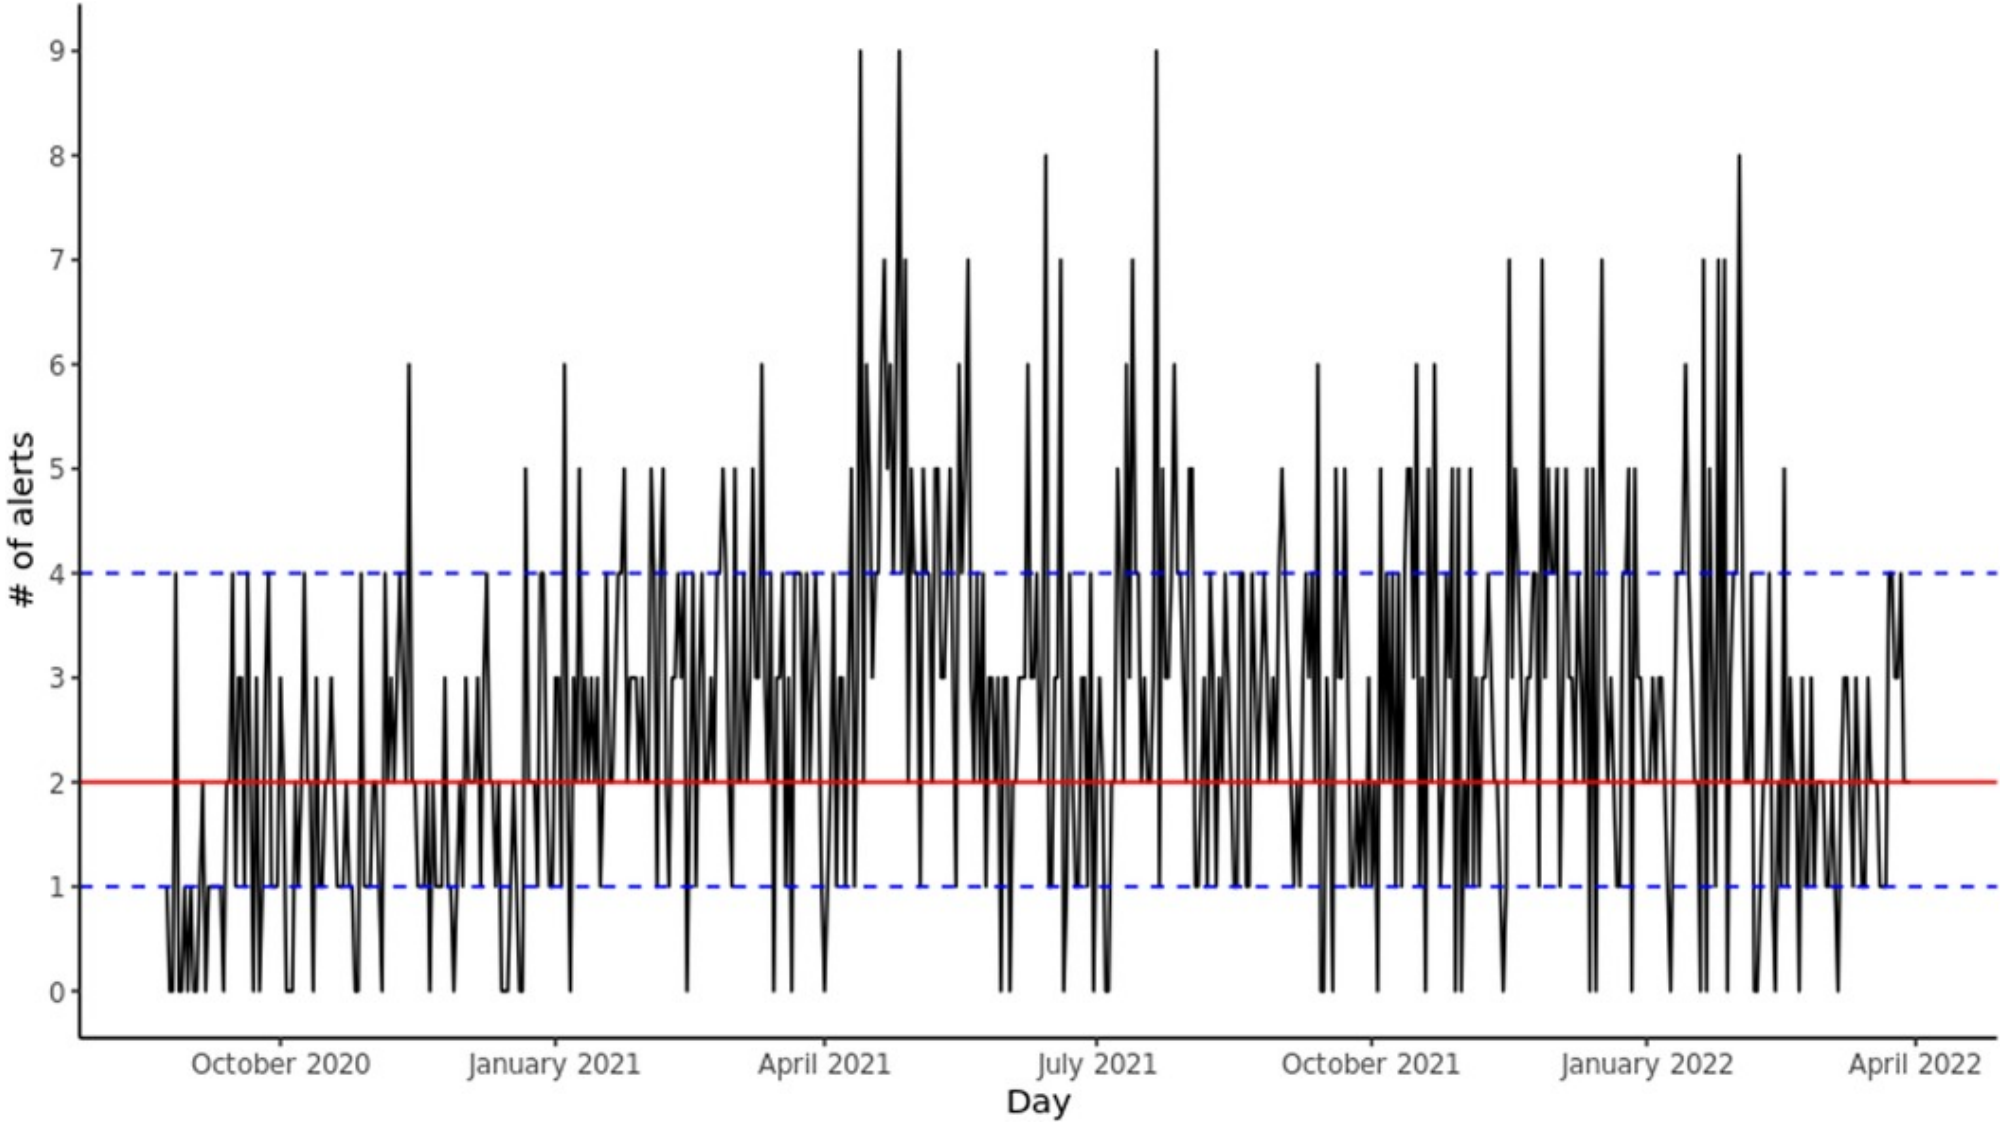

Something about monitoring...?

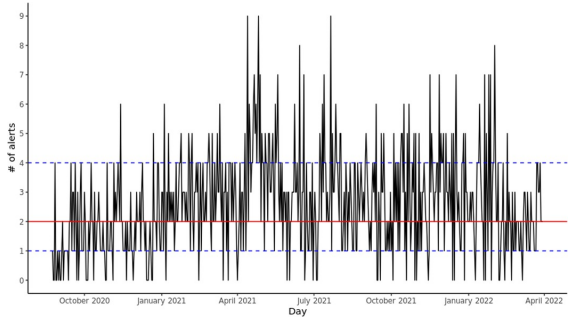

- Minimizing alerts
  - A patient receives an alert if CHARTwatch classifies them as being High risk.
  - Don’t re-alert a patient **unless** it’s been 48 hours since their previous alert.
  - Don’t alert a patient if they just came from the ICU.

Something about monitoring...?

- Clinical adherence
  - # of vitals that are measured after an alert is sent

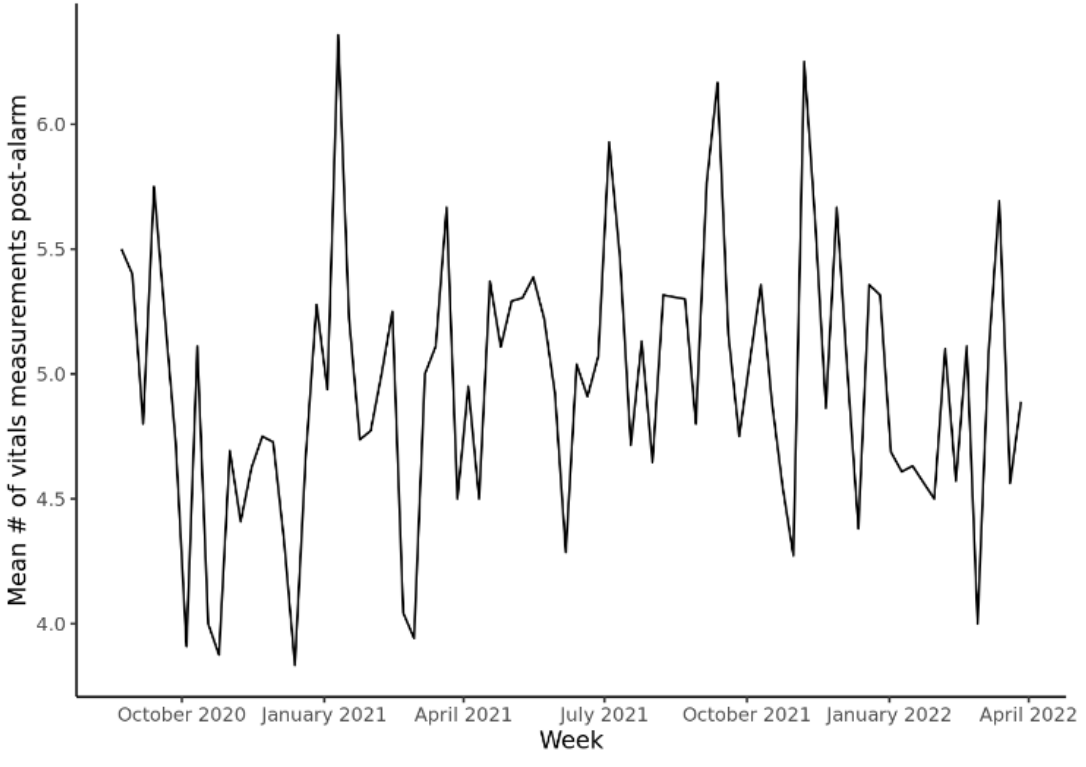

With a little help from my friends

Development:

✓ Database connections

✓ Reproducible environment

✓ Package-based development

✓ `development`, `staging`, `production` environments

Deployment:

✓ Authentication

✓ Downtime protocols

✓ A secure way to download internally-developed packages

✓ Implementation plan

✓ Monitoring

Resources, links, papers, etc.

CHARTwatch

- Papers:
  - Pou-Prom, Chloé, Joshua Murray, Sebnem Kuzulugil, Muhammad Mamdani, and Amol Verma. “From Compute to Care: Lessons Learned from Deploying an Early Warning System into Clinical Practice.” Frontiers in Digital Health, 2022, 174. https://doi.org/10.3389/fdgth.2022.932123.
  - Nestor, Bret, Liam G. McCoy, Amol A. Verma, Chloe Pou-Prom, Joshua Murray, Sebnem Kuzulugil, David Dai, Muhammad Mamdani, Anna Goldenberg, and Marzyeh Ghassemi. “Preparing a Clinical Support Model for Silent Mode in General Internal Medicine.” In Proceedings of the 5th Machine Learning for Healthcare Conference, 950–72. PMLR, 2020. https://proceedings.mlr.press/v126/nestor20a.html.
  - Verma, Amol A., Joshua Murray, Russell Greiner, Joseph Paul Cohen, Kaveh G. Shojania, Marzyeh Ghassemi, Sharon E. Straus, Chloe Pou-Prom, and Muhammad Mamdani. “Implementing Machine Learning in Medicine.” CMAJ 193, no. 34 (August 30, 2021): E1351–57. https://doi.org/10.1503/cmaj.202434.
- Data:
  - Kuzulugil, Sebnem, Chloé Pou-Prom, Muhammad Mamdani, Joshua Murray, Amol A. Verma, Kaiyin Zhu, Michaelia Banning (2022). “GIM, a dataset for predicting patient deterioration in the General Internal Medicine ward (version 1.0.0)”. Health Data Nexus. https://doi.org/10.57764/5rq7-xj70.

DSAA

- Our website: https://unitynet.unity.local/departments-programs-services/corporate-services/data-science-and-advanced-analytics/
- Our blog: https://lks-chart.github.io/blog/
  - “Ohh na na… where are my sodium labs?”

Development tools

- “Introduction to renv” by Kevin Ushey
- “R Packages” by Hadley Wickham and Jenny Bryan

Deployment tools

- ~~RStudio~~ Posit Connect: https://docs.rstudio.com/rsc/
- ~~RStudio~~ Posit Package Manager: https://packagemanager.rstudio.com/client/#/
- Model monitoring:
  - "Model Monitoring with R Markdown" by Julia Silge
  - "MLOps with vetiver in Python and R" by Julia Silge & Isabel Zimmerman
- Feedback loops, "true" false positives and "fake" false positives
  - Adam, George Alexandru, Chun-Hao Kingsley Chang, Benjamin Haibe-Kains, and Anna Goldenberg. 2020. "Hidden Risks of Machine Learning Applied to Healthcare: Unintended Feedback Loops Between Models and Future Data Causing Model Degradation." In *Proceedings of the 5th Machine Learning for Healthcare Conference*, 710–31. PMLR. https://proceedings.mlr.press/v126/adam20a.html.