



EPI Framework Tutorial

Tim Müller (CCI Group, UvA)

t.muller@uva.nl

ICT.OPEN - 20 April 2023

Welcome!

- 12:30 - 12:45: **Introduction** (*presentation*)
- 12:45 - 13:30: **Part 1: Hello, world!** (*guided hands-on*)
- 13:30 - 13:45: **Break**
- 13:45 - 14:15: **Part 2: A workflow for Disaster Tweets** (*hands-on*)
- 14:15 - 14:30: **Evaluation**



Welcome!

- 12:30 - 12:45: Introduction (*presentation*)
- 12:45 - 13:30: **Part 1: Hello, world!** (*guided hands-on*)
- 13:30 - 13:45: **Break**
- 13:45 - 14:15: **Part 2: A workflow for Disaster Tweets** (*hands-on*)
- 14:15 - 14:30: **Evaluation**



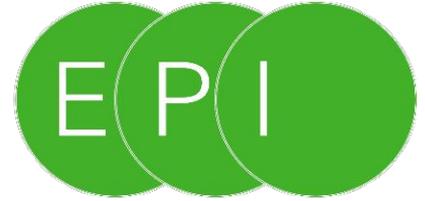
Background

Introducing... Tim Müller

- **23** years old
 - Now: **Scientific Programmer** for the EPI Project
 - Before: Bachelor **Artificial Intelligence** at University of Amsterdam
 - Master **Computer Science, Security Track**, UvA/VU (not completed)
-
- Developer in charge of EPI Framework
 - Combine various parts developed by PhD students
 - Deploy framework at hospitals



Introducing... EPI Project



- EPI -> **Enabling Personalized Interventions**
- Goal: introduction of *Digital Health Twins*
- Various aspects investigated
 - Statistical learning and hypothesis testing
 - Distributed machine learning
 - Privacy-preserving machine learning
 - Automated policy interpretation and enforcement
 - **General-purpose data-sharing framework**

Introducing... EPI Framework

- **EPI Framework** -> general purpose data sharing framework
- Actually, **distributed workflow execution system**
 - i.e., typically, algorithm-to-data
- Focus on **healthcare**
 - Challenges: private datasets -> private policies, autonomous domains
- Built on **BRANE**
 - Focus on separation of concerns and flexibility (hence framework)
- Integrates other technologies
 - BFC Framework
 - Policy reasoners (eFLINT)

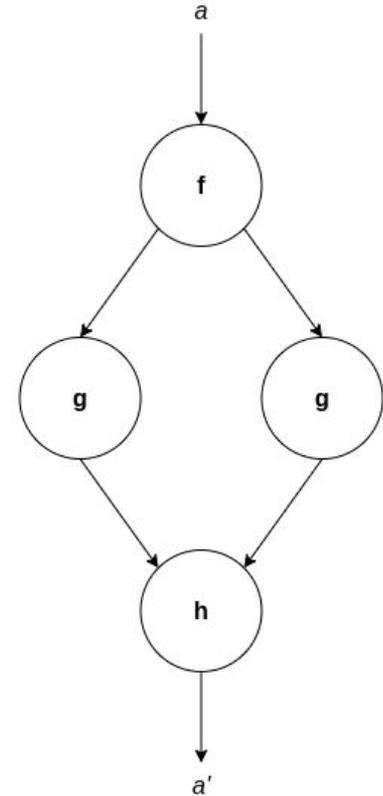
EPI Framework Overview

Workflows

- **High-level, distributed** programs
 - Visualizable as graphs
 - Nodes are **tasks** or **functions**
 - Edges are some **dependency** (typically data)
- Workflow system's job to **fill in details**
 - This act we call *planning*
- Most importantly, **locations** and **ordering** is left to system
 - But the more, the better

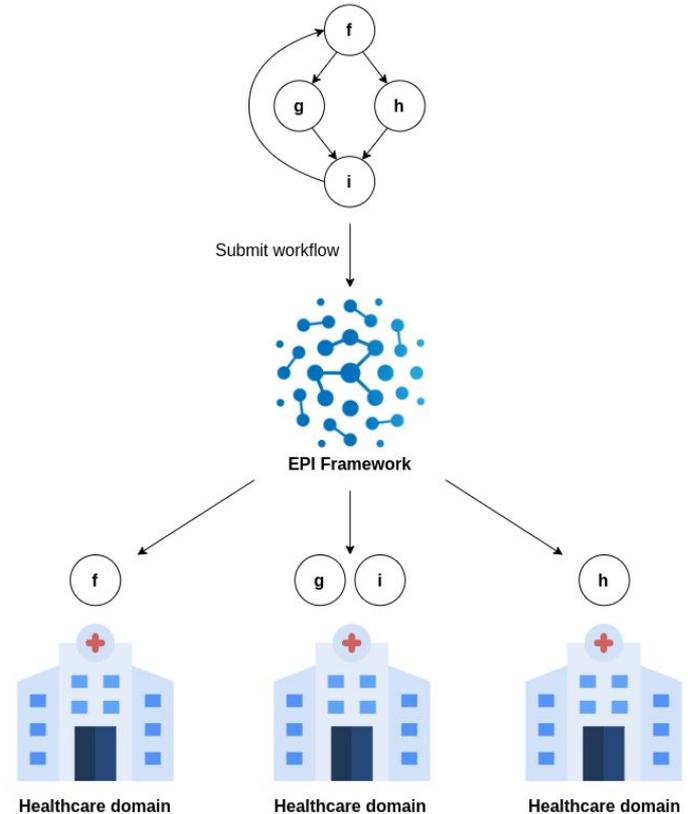


Writes



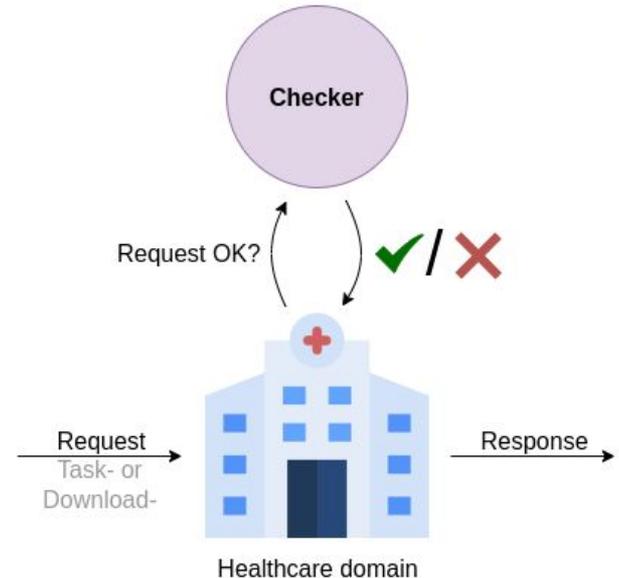
Workflow execution systems

- Typically some kind of **federated system**
- An **orchestrator** is the central part
 - Plans workflows
 - Executes plans
- A **worker** is the local or distributed part
 - Executes individual tasks



EPI Framework

- Orchestrator -> **central node**
- Workers -> **worker nodes**
- Task peculiarities:
 - Tasks are **containerized**
 - Multiple tasks in one container, called a **package**
- Data peculiarities:
 - Workers also host **datasets** or **assets**
 - Workers have **checkers**
 - Checkers enforce **policies**
 - Policies **limit** execution and data access



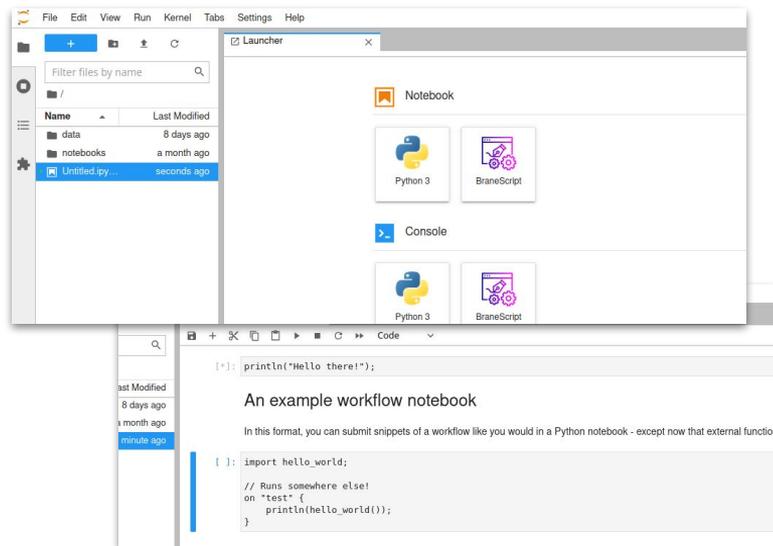
EPI Framework as a user

Client

- A running system is an **instance**
- A **client** is used to submit workflows / upload packages to specific **instance**
- Current client: **brane** executable (CLI tool)
 - JupyterLab integration work-in-progress
- **Local execution** also possible

```
[lut_99@gamelinux brane_docs]$ brane repl
Welcome to the Brane REPL, press Ctrl+D to exit.

1> import hello_world;
2> println(hello_world());
Hello, world!
3> println(hello_world());
Hello, world!
4> println(hello_world());
Hello, world!
5> for (let i := 0; i < 10; i := i + 1) { println(hello_world()); }
Hello, world!
6>
```



Packages

- Role: **Software Engineer**
- Packages are **containerized**
- Language: **any!**
- Metadata specified in **YAML**
- But there is **specific interface**
 - Task selection with **arguments**
 - Input via **environment variables** (JSON)
 - Output via **stdout** (YAML)
 - Data is provided as a **read-only file**

```
1 name: compute
2 version: 1.0.0
3 kind: ecu # Executable Code Unit
4
5 description: |
6 Exposes utilities for preprocessing data, training a classifier,
7 and generating a valid submission file for the Kaggle challenge
8 'Natural Language Processing with Disaster Tweets'.
9
10 contributors:
11 - Andrea Marino <am.marinoandrea@gmail.com>
12 - Jingye Wang <wangjycode@gmail.com>
13
14 dependencies:
15 - python3
16 - python3-yaml
17
18 install:
19 - apt update && apt upgrade -y
20 - apt install pipenv -y
21
22 files:
23 - Pipf --
24 - Pipf 51
25 - _in 52
26 - run_ 53
27 - mode 54
28 - prep 55
29
30 unpack:
31 - pipe 56
32 - pyth 57
33
34 entrypoi
35 kind:
36 exer:
37
38 def main():
39     command = sys.argv[1]
40
41     if command == "create_vectors":
42         # filepath_train_dataset = os.environ["FILEPATH_TRAIN_DATASET"]
43         # filepath_test_dataset = os.environ["FILEPATH_TEST_DATASET"]
44         # filepath_train_vectors = os.environ["FILEPATH_TRAIN_VECTORS"]
45         # filepath_test_vectors = os.environ["FILEPATH_TEST_VECTORS"]
46         # errcode = create_vectors(filepath_train_dataset, filepath_test_dataset,
47                                 #                               filepath_train_vectors, filepath_test_vectors)
48         # print_output({"errcode": errcode})
49
50         # Find the input paths to the training & test dataset
51         train_dataset = f"{json.loads(os.environ['TRAIN_SET'])}/dataset.csv"
52         test_dataset = f"{json.loads(os.environ['TEST_SET'])}/dataset.csv"
53         # Generate the path for the vectors
54         train_vectors = "/result/train_vectors.pickle"
55         test_vectors = "/result/test_vectors.pickle"
56         # Call the function
57         errcode = create_vectors(train_dataset, test_dataset, train_vectors, test_vectors)
58         if errcode != 0: print(f"Uh-oh, 'create_vectors' returned non-zero exit code '{errcode}'", file=sys.stderr)
59
60     return
```

Workflows

- Role: **Scientist**
- Front-end is with **custom Domain-Specific Language (DSL)**
 - **BraneScript**
 - Bakery (WIP)
- Compiles to common Intermediate Representation (IR)
- **Script-like control flow** statements (if, for, while, parallel)
- Can manually specify locations (on-structs)
- Other languages planned:
OpenAPI, CWL

```
3 //
4 // Based on work by Andrea Marino and Jingye Wang.
5 //
6 // A more up-to-date version of the original pipeline defined in 'pipeline'
7 // Specifically, carries the result of functions in the IntermediateResult
8 //
9
10
11 import compute;
12 import visualization;
13
14 on "surf" {
15     ////////// TRAINING //////////
16     println("Cleaning dataset...");
17     let train_clean := clean(new Data{ name := "nlp_train" });
18     let test_clean  := clean(new Data{ name := "nlp_test" });
19
20     println("Tokenizing dataset...");
21     let train := tokenize(train_clean);
22     let test  := tokenize(test_clean);
23
24     println("Removing stopwords from dataset...");
25     train := remove_stopwords(train);
26     test  := remove_stopwords(test);
27
28     println("Performing feature vectorization...");
```

Policies (out-of-scope)

- Role: **Policy Expert**
- Might vary per hospital
- Behind service, so **any language!**
- Possible choices:
 - XACML
 - eFLINT
 - Python
 - Prolog
 - ...
- (Not yet fully implemented)

```
eft players-gamers.eflint
1  // Create the most basic Facts
2  Fact player Identified by String.
3  Fact developer Identified by String.
4  Fact game Identified by String.
5
6  // Create some relational Facts
7  Fact stakeholder Identified by player.
8  Fact created Identified by developer * game.
9  Fact siblings Identified by player1 * player2.
10 Fact can_play Identified by player * game.
11 Fact is_playing Identified by player * game.
12 Fact is_singleplayer Identified by game
13   Holds when (
14     (Exists player : is_playing(player, game)) &&
15     Not(Exists player, player' : player != player' && is_playing(
16   )).
17 Fact is_multiplayer Identified by game
18   Holds when (Exists player, player' :
19     player != player' &&
20     is_playing(player, game) &&
21     is_playing(player', game)
22   ).
23
24 // Create an Event
```

Administration (out-of-scope)

- Role: **System Administrator**
- **branectl** CLI tool
- **Various configuration files**
 - In YAML, mostly

```
[lut_99@workLinux brane_docs]$ branectl -n ../brane/config_test/test
Loading image brane-reg from file ../brane/target/release//brane-reg
Loading image brane-job from file ../brane/target/release//brane-job
Loading image brane-prx from file ../brane/target/release//brane-prx
Running 'docker compose' up on /tmp/docker-compose-worker-2.0.0.yml
[+] Running 4/4
 ✓ Network brane-worker-test Created
 ✓ Container brane-job-test Started
 ✓ Container brane-reg-test Started
 ✓ Container brane-prx-test Started

Successfully launched node of type worker
[lut_99@workLinux brane_docs]$ _
```

```
8 #
9 # For an overview of what you can do in this file, refer to
10 # https://wiki.enablingpersonalizedinterventions.nl/user-guide/system-admins/docs/config
11 #
12
13
14 hostnames:
15   test: 192.168.68.110
16   central: 192.168.68.110
17 node: !worker
18   name: test
19   paths:
20     certs: /home/lut_99/UvA/EPI/BRANE/brane/config_test/test/config/certs
21     packages: /home/lut_99/UvA/EPI/BRANE/brane/config_test/test/config/packages
22     backend: /home/lut_99/UvA/EPI/BRANE/brane/config_test/test/config/backend.yml
23     policies: /home/lut_99/UvA/EPI/BRANE/brane/config_test/test/config/policies.yml
24     proxy: /home/lut_99/UvA/EPI/BRANE/brane/config_test/test/config/proxy.yml
25     data: /home/lut_99/UvA/EPI/BRANE/brane/config_test/test/data
26     results: /home/lut_99/UvA/EPI/BRANE/brane/config_test/test/results
27     temp_data: /home/lut_99/UvA/EPI/BRANE/brane/config_test/test/temp-data
28     temp_results: /home/lut_99/UvA/EPI/BRANE/brane/config_test/test/temp-results
29   services:
30     reg:
31       name: brane-reg-test
32       address: https://brane-reg-test:50055
33       bind: 0.0.0.0:50055
34       external_address: https://test:50055
35     job:
36       name: brane-job-test
```

As a user - Summary

- **Different roles** have **different tools**
 - Hopefully familiar to that role
- Lots of **flexibility** for users
 - Different languages, different interfaces (although not all of them are implemented)
- In this tutorial, you will be a **scientist** and a **software engineer**

Part 1: Hello, world!

Objective

- Write your first **Hello, world!-package**
- See the steps at <https://wiki.enablingpersonalizedinterventions.nl/user-guide>
 - Bottom-left, scroll down to “Tutorials”, then “Part 1: Hello, world!”
 - Or see: <https://tinyurl.com/2xft4cp3>
- I’ll go through it on the board

Tip: Use the jaarbeur hotspot WiFi

Note: The framework is experimental, so unfortunately, expect rough edges

Break

Part 2: A workflow for Disaster Tweets

Objective

- Write a workflow for a package that **classifies disaster tweets**
- See the steps at <https://wiki.enablingpersonalizedinterventions.nl/user-guide>
 - Bottom-left, scroll down to “Tutorials”, then “Part 2: A workflow for Disaster Tweets”
 - Or see: <https://tinyurl.com/mtv24wwp>
- Try to go through the steps yourself!
 - Or just play with the framework :)

Let me know if you have questions or need help!

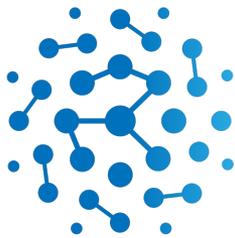
Evaluation

Framework experience?

- Did you find the package design intuitive?
 - Was it too complex?
 - Was the data format (YAML) / language unpleasant?
 - Did running the package locally make sense?
- Did you find workflow writing intuitive?
 - Yes to DSL, no to DSL?
 - DSL intuitive?
 - Is logging in to remote instances clear?
 - Impression about performance?

Tutorial experience?

- Is the material clear enough?
- Were there any unsolvable bugs?
- Did you have enough time?
- Did you like the guided hands-on?
- Did you like the individual hands-on?
- Do you like the topic (scientist/software engineer)?



BRANE

Tim Müller (t.muller@uva.nl)

enablingpersonalizedinterventions.nl

github.com/epi-project/brane

wiki.enablingpersonalizedinterventions.nl (WIP)



The icons (not logos) in this presentation are from Flaticon.com