

Rescue Helicopter Dispatch Analysis and Optimization

Philip Jordan, Christoffer Raun, Xiaoyu Sun, Matus Zilinec
 Hack4Good – Analytics Club at ETH Zürich
 In collaboration with Rega

Abstract—We perform an analysis of past helicopter missions for the Swiss air-rescue Rega. Our main goal is to provide tools that analyze the timing of Rega’s rescue missions and to help with helicopter dispatching in future missions. We predict helicopter flight times for future missions using flight times from past missions, in order to dispatch helicopters more effectively. The dispatching rules are based on expected mission times from the different Rega bases. We make our analysis and predictive models available to Rega dispatchers within an interactive web application.

I. EXPECTED IMPACT

We analyze data about Rega’s historical missions. This may help Rega gain insights into historical missions and to detect outliers (i.e. missions that took a very long time). These insights can then help to optimize timing of future missions, which may improve patient survival rate.

Our predictions of helicopter flight time may help Rega when making dispatching decisions.

Finally, our methods and findings from preprocessing might prove useful in future Rega projects in data collection and analysis.

II. APPROACH

In this section, we will describe our methods for data preprocessing, analysis, and prediction of mission times.

A. Historical data visualization

We visualize historical mission data from Rega’s database. The mission data is first filtered to exclude irrelevant mission types and entries that are too incomplete to visualize. From the available timestamps we then compute the time t_{1a} from emergency call until arrival at the patient, the time t_{1b} spent at the patient’s location, and the flight time t_2 from the patient to the destination hospital. These data are then used to create histogram plots of the time distributions, as well as an interactive map.

B. Preprocessing for flight time prediction

Since the data from Rega’s database is usually collected under time pressure, features are often incorrect or missing. To account for these issues, we utilize multiple preprocessing steps.

First, we discard times of takeoff and landing, which are set by the pilot, often incorrectly. Instead, we use estimate the most likely takeoff and landing time using flight trajectories. We note that the helicopter only sends telemetric data in an interval of 45 seconds, which is then the lower bound of error in any regression tasks.

Second, because of the uncertainty about the exact landing time and place, we split the time taken to reach the patient’s location from the starting point (usually helicopter base), t_1 , in two parts. These are a) t_{1a} , the time taken to reach a 2 km radius around the patient’s location, and b) t_{1b} , the time spent in this radius, including searching for the patient and performing the mission on site. We argue that these two time intervals follow different distributions and depend on different predictors, e.g. a) is mostly influenced by distance and b) might depend on the visibility conditions and terrain at the patient’s location.

Finally, we perform outlier removal. Most importantly, we remove missions planned in advance and missions with multiple flight legs, which are usually required because of special circumstances (e.g. refueling, picking up special equipment).

C. Flight time prediction

In order to predict the expected flight time t_{1a} given a base and a patient location, we use three different approaches, which are described in the following section. The resulting three predictions can then be compared in the dashboard in order to estimate uncertainty.

1) Linear regression model

We apply linear regression to the provided data of past missions using the following features, where $i \in \{1, \dots, k\}$ is the base’s index, $(b_{lat}^{(i)}, b_{lon}^{(i)})$ are the starting GPS coordinates (base location is used when predicting for future missions), (p_{lat}, p_{lon}) are the patient’s coordinates, and $\text{dist}(a, b)$ is the geodesic distance between points a and b .

- 1) $|b_{lat} - p_{lat}|$
- 2) $|b_{lon} - p_{lon}|$
- 3) $\text{dist}((b_{lat}, b_{lon}), (p_{lat}, p_{lon}))$
- 4) $v \in \{0, 1\}^k$ where $v_i = 1$ and $v_j = 0$ for $j \neq i$ (i.e. one-hot encoding of the selected base)

We note that other features, such as elevation, weather, and time of day have also been considered, but didn’t provide any gains in performance. Apart from some of these features being irrelevant, we partially attribute this to low resolution

of the features. The model could also potentially be further improved by extracting text features from the mission type and description.

2) Physical model

The physical model is a simple calculation based on basic physics. Given two points $a = (a_{lat}, a_{lon})$, $b = (b_{lat}, b_{lon})$ and a fixed average flight velocity v_f , the time takes to fly from a to b is

$$t = \frac{\text{dist}(a, b)}{v_f}.$$

In the future, this model could be further improved by considering the acceleration of the helicopter.

3) Nearest neighbor model

In the nearest neighbor model, the expected mission time is predicted as an average over mission times of past incidents that were located close to (p_{lat}, p_{lon}) . If the distance to the nearest neighbor exceeds a certain threshold, we refrain from making a prediction to avoid making inaccurate statements.

When evaluated on historical mission data, the linear regression model outperforms the physical model (a version of which is currently being employed). This comes as no surprise, as in addition to implicitly optimizing over the flight velocity, feature 4) introduces a bias based on the surrounding area, which improves predictions as missions in mountainous regions generally show longer times.

III. RESULTS & DELIVERABLES

Our main deliverable is a JavaScript web application, which can be used to access our results by REGA. This app provides

- interactive visualizations of historical data (Fig. 1),
- heat map of the distribution of average mission times (Fig. 3),
- map of the closest base on average for every location in Switzerland (Fig. 4),
- interface for predicting flight time to an arbitrary patient location using models described in the previous section (Fig. 2). This web application is dynamic, and the user is able to choose the selection of helicopters, mission types, times or positions. The color schemes of both map and plots match and indicate the duration of $t_{1\alpha}$ (where green is short and red is long).

Apart from the web application, we provide a set of Python scripts implementing preprocessing, analysis, and regression on the REGA data.

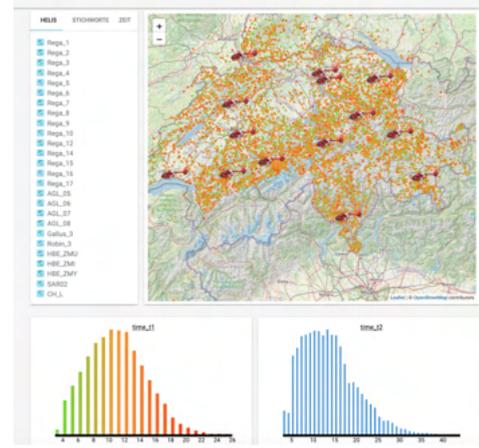


Fig. 1: Historical Data



Fig. 2: Flight time prediction

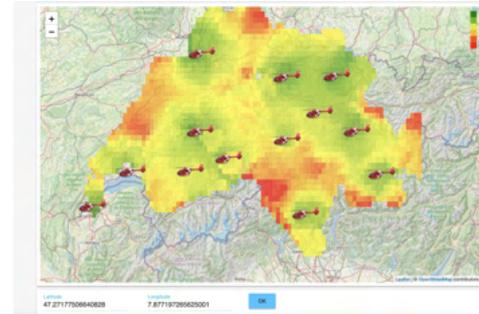


Fig. 3: Helicopter Heat Map

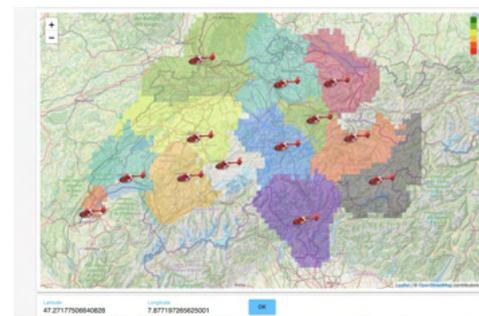


Fig. 4: Service Regions

IV. RECOMMENDATIONS & CONCLUSION

As the interactive dashboard can simply run as an HTML page, the website can simply be set up on a web server and easily accessed by a standard web browser. The historical missions can be analyzed by loading in an excel dump from the current Rega mission control system. The flight time predictions are based on pre-computed models i.e. no further computation has to be performed. One possible improvement could be not to consider the fixed base locations but the current helicopter positions when a dispatch decision needs to be performed. As the helicopter positions are changing constantly, one way would be to automatically provide the dashboard with the current locations of helicopters. A further improvement would be to automatically detect events in the data, e.g. arrival of helicopter at patient's location. This could be done by tracking the helicopter position and would lead to less work for the people in the field and more accurate data.