OPTIMIZING ISSUING POLICIES FOR RED BLOOD CELL UNITS USING DEEP REINFORCEMENT LEARNING

Background

A red blood cell (RBC) transfusion is a common and potentially life-saving medical treatment in which RBCs, originating from a donor, are inserted into the bloodstream of a transfusion recipient to restore the recipient's haemoglobin level and improve oxygen transport. To avoid problems during a transfusion, it is important that *the blood groups* of the donor and recipient *match* (Box 1).

> **Box 1: blood group matching**
>
> The blood group of an individual is determined by the presence or absence of *antigens* on the surface of the RBCs. When a particular antigen is present on the RBCs of the donor, but absent on the RBCs of the recipient, the recipient's immune system may develop *antibodies* against this foreign antigen. These antibodies can cause problems during subsequent transfusions or, in case of a female recipient, during pregnancies. A reaction of the immune system can be prevented by selecting a donor, whose RBCs are *compatible* with the RBCs of the recipient:
>
> - all antigens that are absent on the RBCs of the recipient are also absent on the RBCs of the donor,
> - all antigens that are present on the RBCs of the recipient are either present or absent on the RBCs of the donor.

Hospitals are considering extending their standard matching strategies from ABO, D matching (3 antigens, 8 blood groups) to comprehensive matching (at least 14 antigens, >10,000 blood groups). As a result, Sanquin (the Dutch national blood bank), responsible for the collection and distribution of RBC units, will face an increased demand for comprehensively matched RBC units. A revision of their current inventory issuing policy is necessary within the foreseeable future.

Aim of the project

The aim of the project is

- to develop a deep reinforcement learning (DRL) model for deriving an optimal issuing policy for comprehensively matched RBC units. The issuing policy should minimize the number requests that cannot be satisfied with a comprehensively matched RBC unit (minimize shortages) and minimize the number of RBC units that expire the 35-day lifespan (minimize outdating).
- to add an additional layer on the developed deep reinforcement learning model, which predicts the need for comprehensively matched RBC units of hospitals.

Execution of the project

An optimal issuing policy that finds a trade-off between shortages and outdating can be found by modelling the above problem as a Markov Decision Process (MDP). The drawback of this MDP model is that it cannot be solved within reasonable time, due to the curse of dimensionality (exponentially growing state space). A DRL model does not have this limitation as it determines the next action using a deep neural network instead of enumeration over the entire state space. You will translate the MDP model to a DRL model and determine the best parameter settings.

This project will consist of the following steps:

1. Becoming familiar with the blood group terminology and its mathematical notation.
2. Becoming familiar with the MDP-model.
3. Developing a DRL model for deriving an optimal issuing policy for comprehensively matched RBC units. This DRL model consists of
    a. a deep neural network that determines which RBC unit to select from inventory to satisfy a request for (an) RBC unit(s)
    b. a reinforcement learning module (e.g., Q-learning) to train the neural network
4. Determine the best parameter settings for the deep reinforcement learning model.

Supervision

1st supervisor: Dr. ir. M.P. Janssen
Organization: Sanquin Research & LIACS
E-mail address: m.janssen@sanquin.nl

2nd supervisor: Dr. M. van Leeuwen
Organization: LIACS
E-mail address: m.van.leeuwen@liacs.leidenuniv.nl

Other remarks:

The work will be (part of the time) performed at Sanquin in Amsterdam (days of the week and proportion of time to be discussed with the student). There is an internship fee of €345,- gross per month for this project (excluding cost of travel).