Short-term Modelling of Clinical Variables in Hemodialysis Patients

Michael Thomas Smith, Mauricio A. Álvarez, James Fotheringham
{m.t.smith, mauricio.alvarez, j.fotheringham}@sheffield.ac.uk

The University Of Sheffield.

Background & Motivation

Hemodialysis patients are usually treated three times each week, leading to a gap every week of two days, which is associated with excess hospitalisation and mortality (Fotheringham et al., 2015). Deciding which patients should receive additional treatment could potentially reduce these events. Several studies (e.g. Usvyat et al., 2013) have found measured patient variables are predictive of future hospitalisation and mortality, over a time-frame of days. Rather than focus on predicting hospitalisation we focus on the problem of regressing these variables over a 2-3 day time-window to aid clinical assessment of patient outcomes. To this end we have piloted a system to provide probabilistic predictions of these variables.

This pilot study was conducted on subsets of patients from the 2nd ARO cohort (from 312 EU-FMC facilities between 2007 and 2012). Variables include demographic, dialysis, comorbidity, hospitalisation, mortality and lab results. We focus for this demonstration on blood Potassium (K) and the patient's daily weight gain rate (WGR).

We first demonstrate that lab variables are predictive of hospitalisation, to motivate the study. Figure 1 illustrates both the 2-day break effect (with a greater rate of hospitalisation after the break) and an effect of K also associated with a greater rate of hospitalisation. Unfortunately the lab variables are sampled infrequently. We aim to improve estimates of these variables.

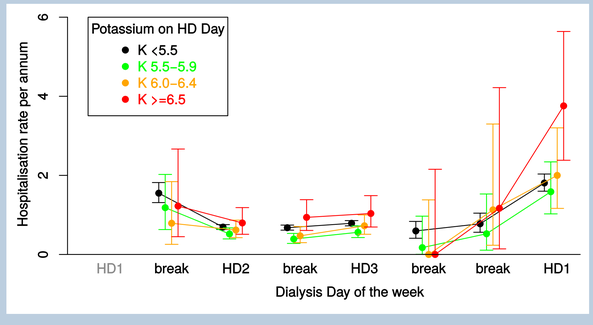

Figure 1: The effect of the dialysis sessions (HD1, HD2 and HD3) on hospitalisation rate. Note the peak after the two day break, and the effect of K on the rate.

Within-patient model

The inputs to the within-patient Gaussian Process (GP) model are vintage (how long the patient has been on dialysis) and the number of days since dialysis (DSD). The outputs consist of variables related to the patient's dialysis (blood flow, weight-change, etc) and lab measurements (e.g. blood Potassium, red blood cell count, etc). An example patient, chosen to demonstrate the correlation, is illustrated in Figure 2.

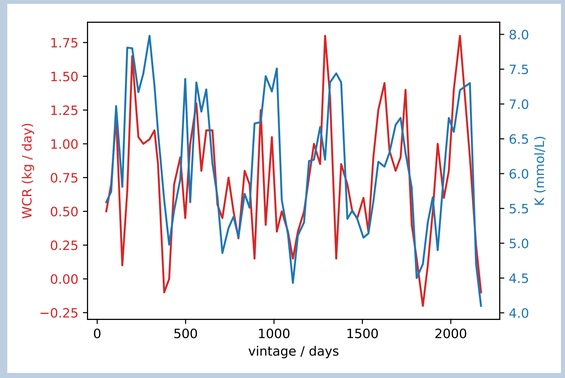

Figure 2: Patient 3118's K and WCR over six years.

We confirmed that the GP performs more accurately than simply using the mean of the prior three measurements. We chose to assess the accuracy of predictions for the Weight Change Rate (weight increase per day between dialysis sessions, WCR) and the lab blood potassium (K) values, for a randomly selected group of 158 patients (with 25 lab sessions or more). We fitted a GP with an exponentiated quadratic kernel (EQ) to each patient's data, optimising the hyperparameters. We found the RMSE errors were reduced by 7.6% and 10.4% for the WCR and K respectively ($p<0.05$, determined using bootstrap).

We next investigated the potential benefits of coregionalisation. We found, belatedly, that correlation structures within individuals are very hetrogenous across the population. We found that variables we expected to be strongly correlated, were only correlated in a subset of patients. Figure 3 demonstrates the variation in correlation between the WCR and K (for patients with >20 lab values). We tested coregionalisation with a tiny subset of 13 patients with the greatest correlation ($p<0.000001$).

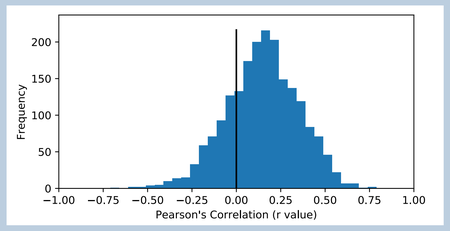

Figure 3: Distribution of correlations amongst patients between K and WGR.

We found a significant improvement in the RMSE of the predictions of K over the non-coregionalised model (0.44 from 0.52 mmol/L, $p<0.05$) but not for the more frequently measured WGR. This significance was quite fragile and was weakened by including patients with less correlation.

Prior model

Towards the beginning of a patient's time series, we have very little data to estimate the prior mean. We thus investigated using an estimate from the population. For this example we selected 500 patients to train the population prior model with. We chose five simple demographic features (age, gender, vintage, height and initial weight) and used these to predict the WCR and K. We then tested this by looking at predictions from a different set of 1000 patients. In which we compare the prediction error when using the mean provided by the population model. We did not find a robust reduction in RMSE. Figure 4 illustrates one axis of the population model, showing how WGR varies with age, but also how the scale of the noise is considerably greater than the signal.

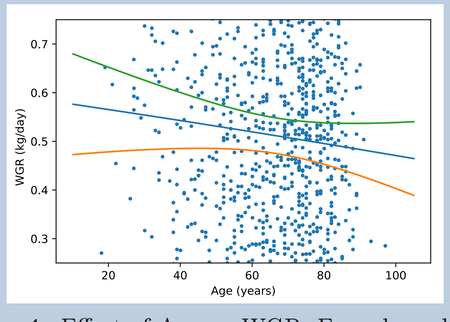

Figure 4: Effect of Age on WGR. Error bars describe latent function 95% CI.

Hierarchical Model

One hypothesis was that rare but important correlated behaviour exists between variables that an individual, coregionalised model won't be able to learn due to the rarity of the events. For example; a splike in white blood cells is commonly associated, anecdotally, with a drop in weight. We experimented with a hierarchical model in which the errors in the individual model are used as outputs in another 'population' model. As inputs we experimented using features such as the rate of change or absolute difference in variables. To predict we added estimates from this population model to the predictions from the individual model. We unfortunately found little effect (at most a 2% decrease in the RMSE for the WCR was detected). We found that, using aggregated statistics, the correlations, although statistically significant, represented largely undetectable improvements at the level of an individual.

Conclusions and Future Work

These experiments have demonstrated the challenges involved in using clinical data for short-term forecasting. In future we will test the uncertainty estimates. We have already confirmed the model's 95% CIs enclose about 95% of the predictions. However, in the run-up to hospitalisation the proportion may change. Work is required to see whether the hierarchical model can improve this. Finally, it is interesting to note that the presence of correlation does not mean that coregionalisation will aid prediction. The derivative correlations may be more relevant.

This work has been supported by the Engineering and Physical Research Council (EPSRC) Research Project EP/N014162/1.

Fotheringham J, Fogarty D G, El Nahas M, Campbell M J and Farrington K 2015 *Kidney international* **88**(3), 569–575.

Usvyat L A, Kooman J P, van der Sande F M, Wang Y, Maddux F W, Levin N W and Kotanko P 2013 *Nephrology Dialysis Transplantation* **29**(2), 442–448.