Gaussian Processes for Prediction in Intensive Care

Fabián Güiza FABIAN.GUIZA@CS.KULEUVEN.BE

Jan Ramon JAN.RAMON@CS.KULEUVEN.BE

Hendrik Blockeel HENDRIK.BLOCKEEL@CS.KULEUVEN.BE

Department of Computer Science, K.U.Leuven, Celestijnenlaan 200A, B-3001 Heverlee, Belgium

Abstract

In this paper we present the use of Gaussian Processes for regression in the application of prediction in Intensive Care. We propose a preliminary solution to predicting the evolution of a patient's state during his stay in intensive care by means of defined patient specific characteristics.

1. Introduction

In intensive care, automatic online measuring systems register and store over 200 physiological and device parameters that together describe the state of a critically ill patient (Gather et al., 2000). The resolution of this information can be very high with parameters being registered every minute. It is crucial in intensive care to detect clinical problems early enough so that preventive or curative treatments can be applied in time.

In practice, a physician analyses all the patient related data in order to foresee a change in the patient's condition and administer the appropriate treatment. Humans can not deal well with problems involving more than seven variables, and are not able to judge degrees of relatedness for more than two variables. Consequently, of all the available information a physician usually selects only a few variables that he considers important (based on his experience) to assist him in his decision-making process.

It is of interest to develop reliable procedures that automatically analyse typical intensive care data, and that are well equipped to cope with its dynamical multivariate nature. These methods should provide reliable predictions (with an associated confidence value) of the patient's future states so that the physician may have a clearer picture of the patient's evolution and can apply the appropriate treatments.

2. Application

We are interested in developing models that use the data generated for patients in intensive care to predict future values of variables that are considered interesting by physicians. These variables are considered important in determining the future state of a patient and therefore his evolution in an Intensive Care Unit (ICU).

Patients have individual characteristics, which are assumed to be constant for a given patient but different amongst individual patients, e.g. patients have individual reaction times to certain medications. These characteristics play an important role since they define the 'normal' or 'typical' state of a patient. Deviations in the monitored measurements with respect to these patient specific characteristics is valuable information when predicting a change in the state of the patient. For example a large deviation in a patient's creatinine level with respect to his 'normal' level upon admission in ICU, is a clear indication of an acute failure in the functions performed by the patient's kidneys.

The inclusion of these individual patient characteristics as inputs for predictive models would clearly have a great beneficial effect on their predictive performance. These characteristics are in general not known upon admission of a patient to ICU, but can be estimated from the daily measurements, with estimates converging to the real values after sufficient observations. Even though upon admission to the ICU the estimates of the patient's specific characteristics will be uncertain, we can use the information of previously observed patients to define a prior over these characteristics, thus bounding the initial uncertainty. For example after having observed a population of patients during several days, the average heart-rate for each patient can be computed with low uncertainty, and a distribution of average heart-rate over all patients can be determined. This distribution will be the prior for a new patient's average heart-rate. Given the uncertainty related to the specific patient characteristics, the models we are interested in developing should allow for predictions with uncertain inputs.

Gaussian Processes for regression have been used to model and forecast real dynamic systems, because of their flexible modelling abilities and their high predictive performances. They allow for multi-dimensional inputs, and they assign a confidence value (variance) to their predictions (Rasmussen & Williams, 2006). These properties make Gaussian Process an appropriate modelling tool for our application. Gaussian Processes have recently been extended to include predictions on noisy or uncertain inputs (Girard, 2004), thus allowing the use of the uncertainties related to the patient specific characteristics. Not taking this uncertainties into account will result in a model that makes over-confident predictions, and this can be potentially dangerous in the current application where the predictions are used for critical decision making processes on the physician's part.

3. Methods

In the following we will give a brief overview of the key computations required when training a Gaussian Process for regression and how the trained model can be used for prediction on new inputs, with special consideration of noisy or uncertain inputs. A detailed description of Gaussian Process in general can be found in (Rasmussen & Williams, 2006) and of their application to noisy inputs in (Girard, 2004). We end this section with a discussion of how the specific patient characteristics are updated as observations occur for a specific patient during his stay in intensive care.

Given a data set $\mathbf{D}$ of D-dimensional input vectors $\mathbf{x}_i$ and their corresponding observed outputs y_i, which are related by

$$y_i = f(\mathbf{x}_i) + \epsilon \tag{1}$$

where $\epsilon \sim \mathcal{N}(0, \sigma_n^2)$ is additive Gaussian noise. The nonlinear mapping $f(.)$ is modelled by a zero-mean Gaussian Process so that

$$p(\mathbf{f}|\mathbf{D}) \sim \mathcal{N}(0, \mathbf{K}) \tag{2}$$

where the covariance matrix $\mathbf{K}$ of the process is computed from the covariance function or kernel and $\mathbf{K}_{ij} = k(\mathbf{x}_i, \mathbf{x}_j)$. From the previous equations it can be seen that

$$p(\mathbf{y}|\theta, \mathbf{D}) \sim \mathcal{N}(0, \mathbf{K} + \sigma_\epsilon^2 \mathbf{I}) \tag{3}$$

where $\mathbf{I}$ is the identity matrix and θ is the vector of the parameters of the covariance function and the noise variance, which together constitute the model's hyper-parameters. When training the model a maximum likelihood approach is followed to find values for the hyper-parameters such that the evidence is maximised.

A common choice for a covariance function is the squared exponential or Gaussian kernel

$$k(\mathbf{x}_p, \mathbf{x}_q) = \sigma_f^2 \exp(-\frac{1}{2}(\mathbf{x}_p - \mathbf{x}_q)^T M (\mathbf{x}_p - \mathbf{x}_q)) \tag{4}$$

for $M = \mathrm{diag}(m_1, ..., m_D)^{-2}$ and each m_d is a characteristic length scale for each input dimension. For this choice of covariance function the hyper-parameters are σ_f^2 and the m_d characteristic length scales. The values found for the m_d when training the model determine the relevance of the corresponding input dimension, in the so called Automatic Relevance Determination (ARD).

The trained model can now be used to determine the predictive distribution of $f(\mathbf{x}_*)$ at an input $\mathbf{x}_*$.

$$
\begin{aligned}
p(f(\mathbf{x}_*)|\mathbf{x}_*,\mathbf{D}) &\sim \mathcal{N}\left(\mu(\mathbf{x}_*),\sigma^2(\mathbf{x}_*)\right)\\
\mu(\mathbf{x}_*) &= \mathbf{k}(\mathbf{x}_*)[\mathbf{K}+\sigma_n^2 I]^{-1}\mathbf{y}\\
\sigma^2(\mathbf{x}_*) &= k(\mathbf{x}_*,\mathbf{x}_*)-\mathbf{k}(\mathbf{x}_*)^T[\mathbf{K}+\sigma_n^2 I]^{-1}\mathbf{k}(\mathbf{x}_*)
\end{aligned}
\tag{5}
$$

In our application we are interested in a prediction on an input $\mathbf{x}$ (corresponding to patient specific characteristics) that has an associated uncertainty, $\mathbf{x} = \mathbf{u} + \epsilon_x$ and $\epsilon_x \sim \mathcal{N}(0, \mathbf{S})$ or equivalently $\mathbf{x} \sim \mathcal{N}(\mathbf{u}, \mathbf{S})$, the predictive distribution must be integrated over the input distribution

$$
p(f_*|\mathbf{u},\mathbf{S},\mathbf{D}) = \int p(f_*|\mathbf{x}_*,\mathbf{D})p(\mathbf{x}_*|\mathbf{u},\mathbf{S})d\mathbf{x}_* \tag{6}
$$

However, since $p(f_*|\mathbf{x}_*,\mathbf{D})$ is a nonlinear function of $\mathbf{x}_*$ then $p(f_*|\mathbf{u},\mathbf{S},\mathbf{D})$ is not Gaussian and the integral must be solved by performing some approximations.

A first possible approximation to the true distribution results of using numerical methods that rely on Markov-Chain Monte-Carlo sampling

$$
\begin{aligned}
p(f_*|\mathbf{u},\mathbf{S},\mathbf{D}) &= \int p(f_*|\mathbf{x}_*,\mathbf{D})p(\mathbf{x}_*|\mathbf{u},\mathbf{S})d\mathbf{x}_*\\
&\simeq \frac{1}{T}\sum_{t=1}^{T} p(f_*|\mathbf{x}_*^t,\mathbf{D})
\end{aligned}
\tag{7}
$$

and $\mathbf{x}_*^t$ are samples from $p(\mathbf{x}_*|\mathbf{u},\mathbf{S})$

An alternative approximation is the analytical Gaussian approximation in which the first two moments (mean and variance) of $p(f_*|\mathbf{u},\mathbf{S},\mathbf{D})$ are computed, so that

$$
p(f_*|\mathbf{u},\mathbf{S},\mathbf{D}) \simeq \mathcal{N}(m(\mathbf{u},\mathbf{S}), v(\mathbf{u},\mathbf{S})) \tag{8}
$$

Depending on the type of Covariance function used, these moments can be determined exactly or via some approximation. Exact solutions exist for a kernel which is Gaussian or linear. Equations for computing these moments are given in ((Girard, 2004))

The patient specific characteristics used here as uncertain inputs for prediction can be modelled as Gaussian distributions, which are computed by looking at large patient populations. As was previously discussed, these distributions serve as priors ($\mathbf{x} \sim \mathcal{N}(\mathbf{u}_0, \mathbf{S}_0)$) for a Bayesian Inference process carried out for each new patient under observation. After each time step, the patient specific characteristics are recomputed by using the current observation to update the prior. The mean of the obtained posterior will thus be closer to the real value and the uncertainty will decrease. After sufficient observations have occurred, the variance associated with the patient's specific characteristic will be small resulting in an overall decrease in the variance of the model's predictions.

4. Future Work and Challenges

A first step towards developing the work outlined here lies in defining a good set of patient specific characteristics that can be considered relevant in the different prediction tasks of interest. Some of these characteristics will be known to domain experts and will be easily incorporated into the models. However for some prediction tasks, parameters that can be used as patient specific remain to be discovered, leading to the interesting task of learning them from the data. As an example, patients could be clustered using a similarity measure on one of their time varying parameters, and a patient specific characteristic would be the cluster to which he belongs.

Training of Gaussian Process models with the selected set of inputs and making use of Gaussian kernels will allow to determine which patient specific characteristics have a real impact in predicting the desired variables, a knowledge discovery resulting of Automatic Relevance Determination. However, given the large amount of characteristics that can be defined as patient specific, other techniques might be required in order to search through this space for the appropriate characteristics for a given prediction task.

The use of real data to experimentally evaluate the predictive performance of the Gaussian Processes will allow to verify that true function values remain within computed confidence intervals. This will determine if the models developed are not over-confident, and therefore if their predictions can be used as input to decision-making processes.

The predictive performance of models built with different kernel functions can be compared for the tasks of predicting different time-series,e.g heart-rate, temperature. Since there will be uncertainty associated to the inputs, the computation of some of these predictions will require the use of MCMC methods. This is the case of Neural Network kernels which might prove valuable when predicting time-series that fall outside the smoothness assumption of Gaussian kernels.

Extensions to Gaussian Process framework that allow for training with noisy inputs (Girard, 2004) are also interesting to consider since the recorded measurements used for training will typically be corrupted by some level of noise (Gather et al., 2000).

It must be mentioned that because of the large amount of data that is available through the use of online monitoring, sparse methods will also have to be considered to avoid the computation of large Kernel matrices.

5. Conclusion

We have presented a preliminary solution to the problem of predicting the evolution of a patient's state during his stay in ICU by means of patient specific characteristics and Gaussian Process that deal with their inherent uncertainties. As discussed in the previous section, a number of challenging aspects remain to be addressed in order for the solution here described to be applicable. Applying this solution to real data is considered as future work.

6. Acknowledgements

Fabián Güiza is supported by the Interdisciplinary Research Fund of the Katholieke Universiteit Leuven (IDO 3M030483). Jan Ramon and Hendrik Blockeel are post-doctoral fellows of the Fund for Scientific Research (FWO) of Flanders.

References

Gather, U., Imho, M., & Fried, R. (2000). *Graphical models for multivariate time series from intensive care monitoring* (Technical Report 33/2000). University of Dortmund, Dortmund, Germany.

Girard, A. (2004). *Approximate methods for propagation of uncertainty with gaussian process models.* Doctoral dissertation, Department of Computing Science, University of Glasgow.

Rasmussen, C., & Williams, C. (2006). *Gaussian processes for machine learning.* Cambridge, Massachusetts: MIT Press.