A Database of Mood Time Series

P.J. Moore[1] P. McSharry[2] M. Little[3]

[1]Oxford Centre for Industrial and Applied Mathematics
University of Oxford

[2]*Supervisor.* Smith School of Enterprise and the Environment
University of Oxford

[3]*Supervisor.* MIT Media Lab, Cambridge, Massachusetts

Trinity Term 2011

Talk Outline

- Introduction
  - Bipolar disorder, the database and how the data was collected
- Data characteristics and cleaning
  - Summary statistics and cleaning criteria
- Forecasting
  - Using standard time series methods and Gaussian process regression
- Effects of gender, season and environment on mood
- Further work
  -Models of mood dynamics, clustering by dynamics, controlling mood

Bipolar Disorder

People with bipolar disorder suffer extremes of depression and mania. Depression can mean persistent unhappiness, loss of appetite and problems with sleep, while mania sometimes leads to poor choices: gambling, spending sprees and sexual indiscretions. The disorder often affects creative people, among them Sinead O'Connor, Stephen Fry and - bearing in mind the issues of retrospective diagnosis - the mathematician Georg Cantor.

Measuring Depression

Inventory Category	*Domain Score (0-3)*
1. Sleep *(4 questions)*	
2. Feeling sad	
3. Appetite/weight *(4 questions)*	
4. Concentration	
5. Self-view	
6. Death/suicide	
7. General interest	
8. Energy level	
9. Slowed down/Restless *(2 questions)*	

The QIDS Scale for depression: domains and symptoms. There is more than one question for domains 1, 3 and 9, so there are 16 questions in total. The score in these cases is calculated by taking the maximum score over all questions in the domain. The total score is the sum of the domain scores.

Measuring Mania

Inventory Category	*Item Score (0-4)*
1. Feeling happier or more cheerful than usual	
2. Feeling more self-confident than usual	
3. Needing less sleep than usual	
4. Talking more than usual	
5. Being more active than usual	

The Altman self-rating mania scale. Each question can score from $0-4$ giving a maximum possible score of 20.

Monitoring Mood in Bipolar disorder

A mood monitoring system

Since 2006, the Department of Psychiatry in Oxford has been monitoring mood in patients with bipolar disorder by using SMS text messages to return questionnaire results. Each week, participants return two numbers representing the depression and mania scores. The resulting database now holds mood ratings for more than 400 patients over 5 years.

Example Depression Time Series

Examples of depression time series

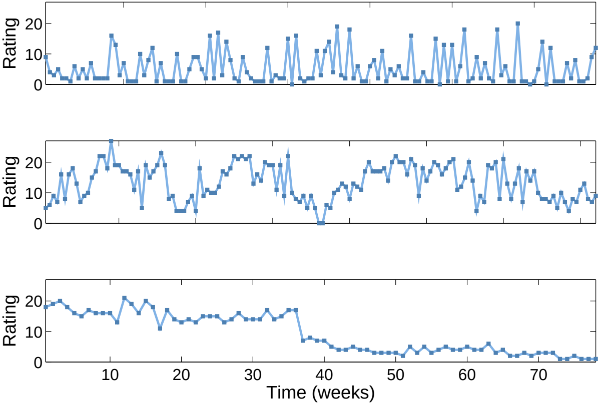

Example Mania Time Series

Examples of mania time series

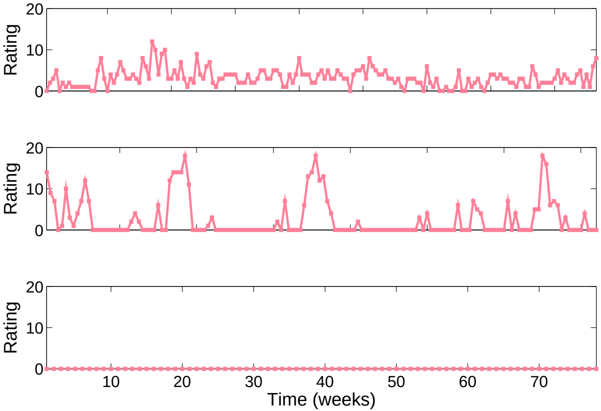

Time Series Lengths

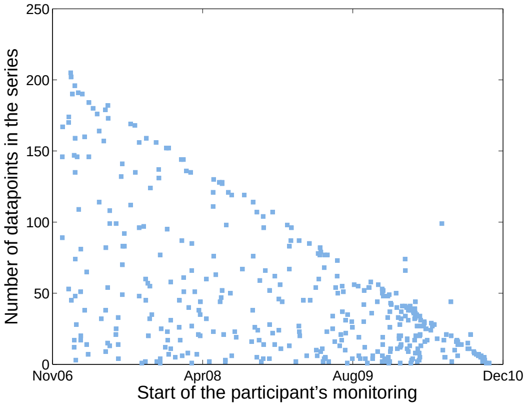

Figure: Number of data point in each patient's time series. The maximum length usually depends on when the patient joined the monitoring scheme. The outliers are patients who return data more frequently than every week.

Histogram of Depression Ratings

Histogram of Depression Ratings

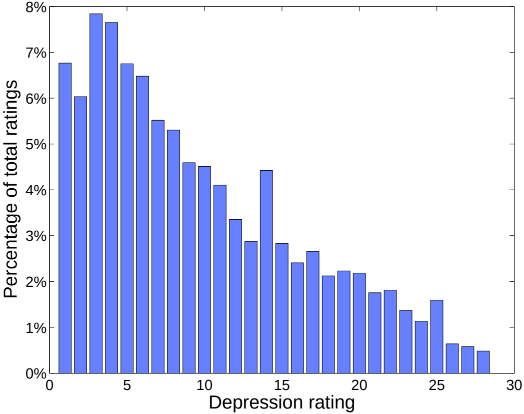

Figure: Histogram over returned depression ratings pooled over all patients.

Histogram of Mania Ratings

Histogram of Mania Ratings

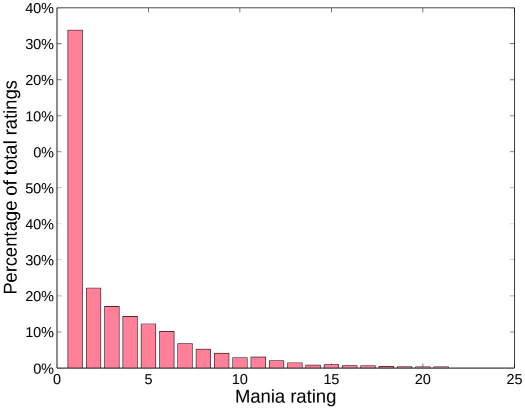

Figure: Histogram over returned mania ratings pooled over all patients.

Proportion Of Zeros

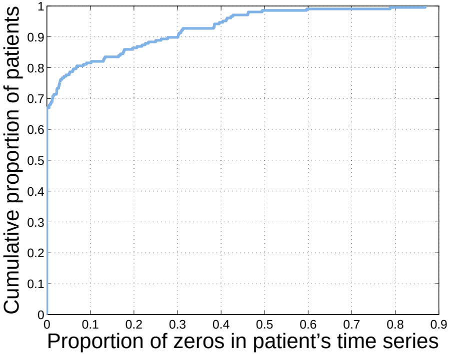

Figure: Empirical cumulative probability density of the proportion of zeros in the depression time series of patients with a least 30 data points.

Response Intervals

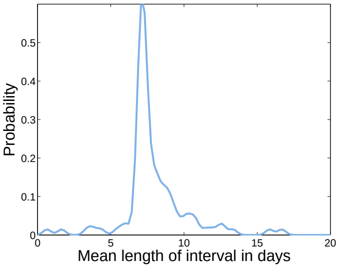

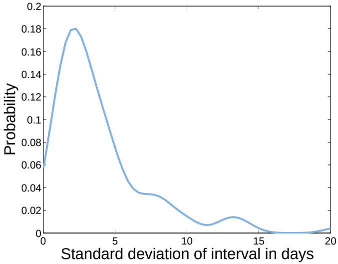

Mean and standard deviation of the time difference between observations. The distributions are estimated using a Gaussian kernel at evenly spaced 100 points.

Selection Criteria

Filter criterion	*Condition (units in days)*	*Failed*
Minimum observations	$N_{obs} \geq 30$	186
Maximum zero ratings	$N_{zeros} \leq 0.5 N_{obs}$	4
Mean/SD of interval	$5 \leq \mu \leq 14, \sigma \leq 10$	99

Filter criteria for selecting the data set for forecasting. A given time series may fail more than one condition, so for example many short series with fewer than 30 points also fail the condition for the response statistics. Starting with 397 participants, 170 are left after the filtering process is complete.

Non Stationarity of Time Series

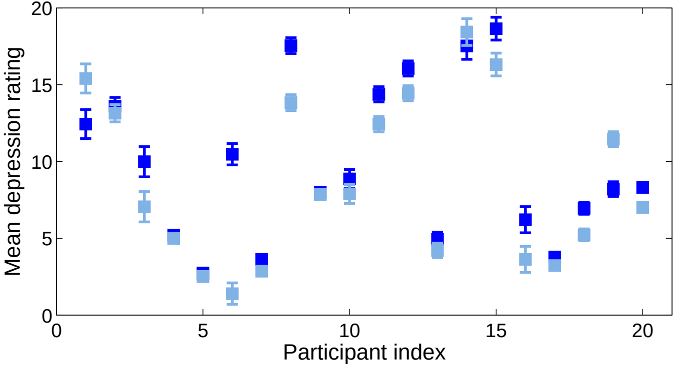

Dark blue markers represent the sample mean of the first half and light blue the second half. The error bars represent the standard error assuming that the numbers are Gaussian distributed and uncorrelated.

Scaling Exponents

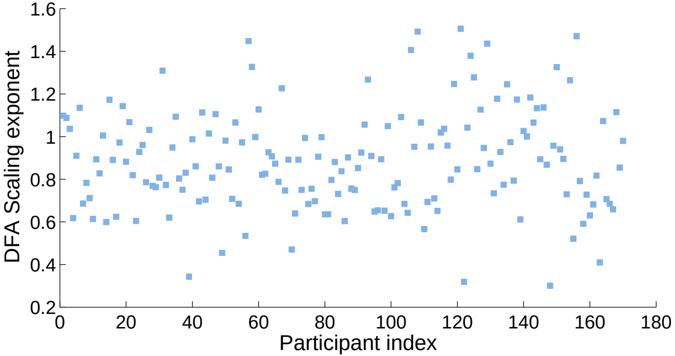

The scaling exponent is derived from detrended fluctation analysis (DFA) for each of the time series in the cleaned data set. It can be seen that the data set is heterogeneous.

Issues in forecasting

- What is the *purpose* of forecasting?
  - Clinical utility - but we are using a cleaned subset of the data
  - *Spin offs*: we can examine a succesful model to see what can be learned about mood dynamics
- How do we *train* the model parameters? Over all patients or per patient
- Which forecasting *methods* do we use?
  - Simple exponential smoothing
  - Gaussian process regression

Exponential Smoothing

The simplest form of exponential smoothing takes a previous forecast and adjusts it with a forecast error. If we denote the time series value as y_t, the existing forecast of that value as $\hat{y}_t$, then the forecast for the next period is

$$\hat{y}_{t+1} = \hat{y}_t + \alpha(y_t - \hat{y}_t) \tag{1}$$

where α is a constant between 0 and 1. Hence

$$\begin{aligned} \alpha = 1 &\Rightarrow \hat{y}_{t+1} = y_t \\ \alpha = 0 &\Rightarrow \hat{y}_{t+1} = \hat{y}_t \end{aligned} \tag{2}$$

Simple exponential smoothing is equivalent to an ARIMA(0,1,1) model, that is, a differenced first order moving average model with no constant.

Parameter estimation

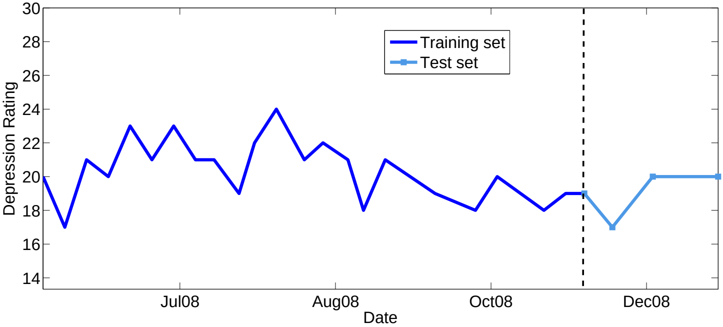

- We define $\hat{y}_1$ to be the mean of the training set.
- The exponential smoothing parameter $\hat{\alpha}$ is chosen to minimise the RMSE error over the training window.

Parameter Stability

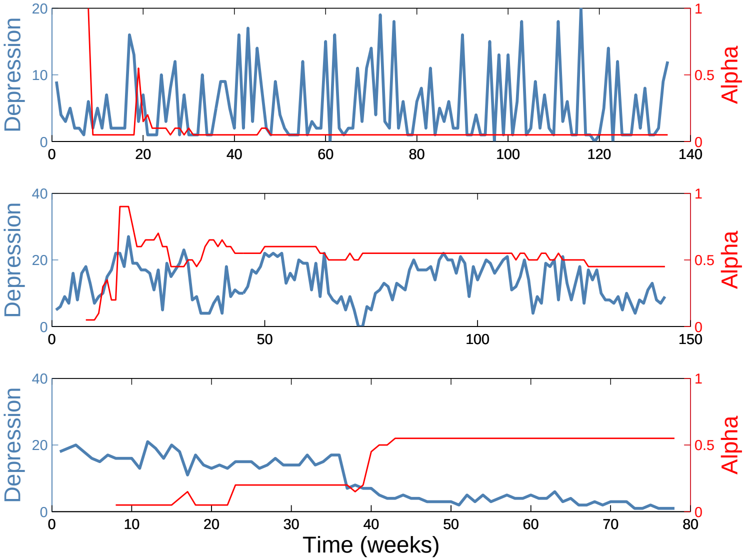

Parameter Stability - 2

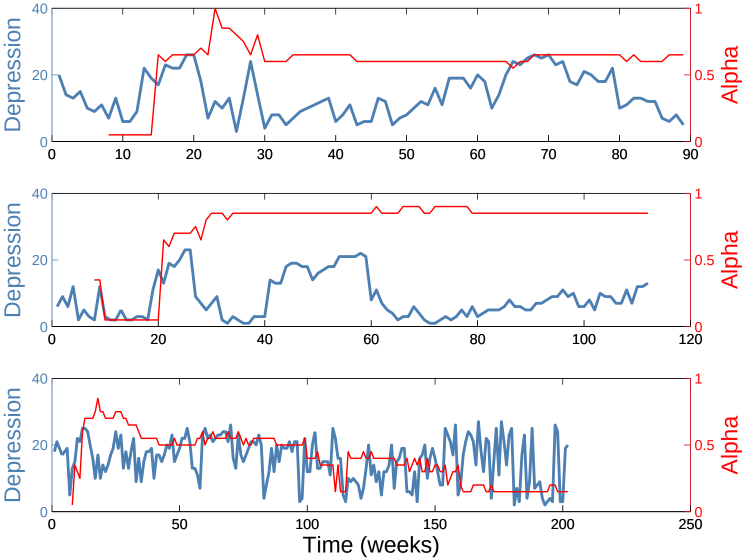

Gaussian Process Regression

- For a Gaussian Process the joint distribution between any finite number of these variables is Gaussian. The process is specified by a mean function and a covariance function.
- We postulate a certain form of parameterised covariance function and estimate the parameters from the data.
- The expected value and variance of the function can then be found for any input.

Gaussian Process Forecasting

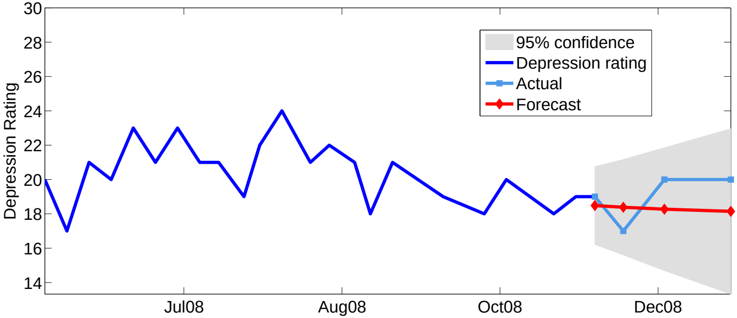

Figure: Forecasting depression in bipolar disorder. The depression ratings (dark blue line) are used to make forecasts (red triangles) of future ratings. The actual measured values are shown in light blue.

Forecasting Approach

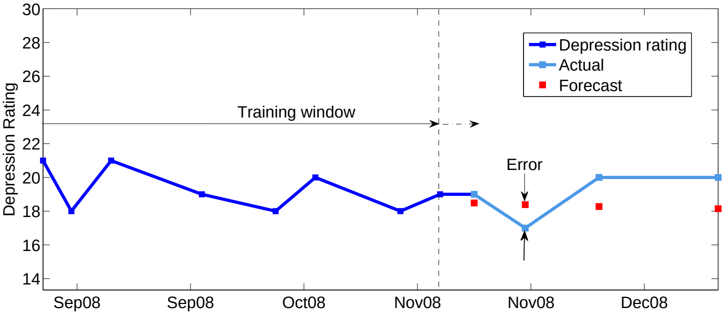

For each time series, a window is extended stepwise from a minimum offset to the end of the time series. At each step, a forecast is made for the next data point in time outside the window, and the difference between the forecast and the true value is recorded. The root mean square error (RMSE) for the test estimations is then computed and used as a summary error for the participant.

Forecast Results

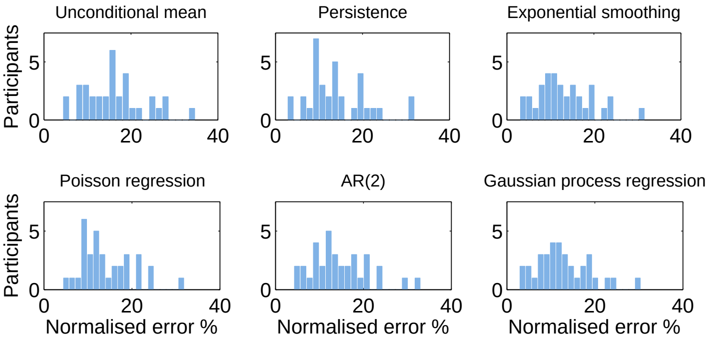

Histograms of errors for a test set of 34 participants using an unpartitioned training set. These results are percentages obtained by dividing by the range of the rating scale (27) and multiplying by 100.

Psychological Interpretation

Psychological interpretation in the light of self esteem modelling:

- Self esteem has been found to be modelled using exponential smoothing (Fortes, *Dynamics of Self-Esteem and Physical Self*).
- They suggest that two processes underlie the dynamics of self-esteem: *preservation* which tends to restore the previous value after a disturbance and adaptation, which tends to inflect the series in the direction of the perturbation
- The disturbance (or shock) can be seen as the result of all the good and bad events (known in psychology as *stressors*). The value of α might then provide a measure of the preservation/adaptation spectrum.

Gender Differences

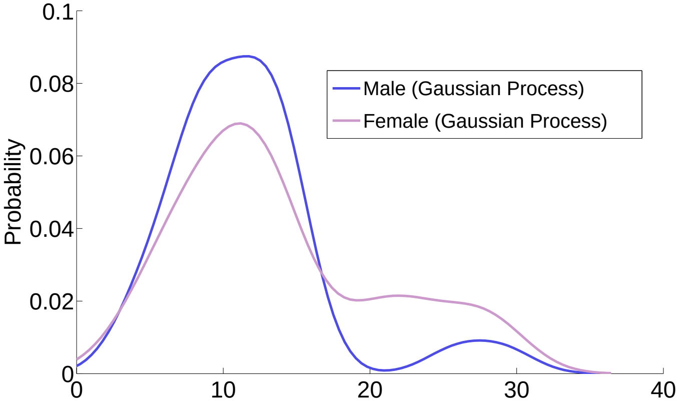

Figure: Kernel density estimation for distribution of Gaussian process forecast errors for two sets of 30 patients, one all male and the other all female ($p = 0.02$). The x-axis is the RMS forecast error for a participant over a test set, normalised by the maximum value of the rating scale.

Seasonal Effects

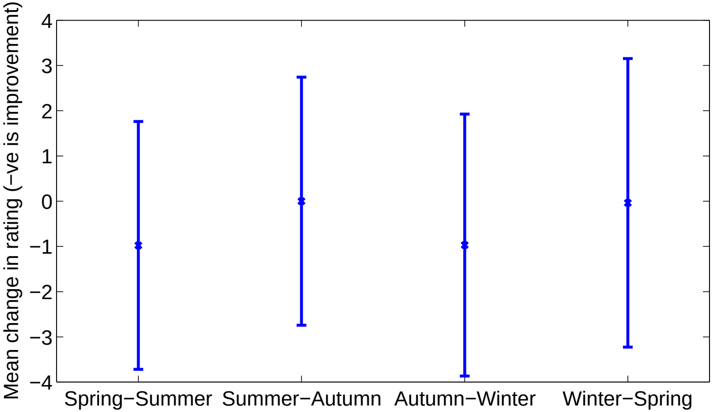

Figure: Mean change in depression over the cleaned data set by season. Standard errors are shown, assuming a Gaussian distribution.

Environmental Effects

ID	Tmin	Tmax	Rain	Sun	Humidity	Wind speed	TCI
1	0	0	0	0	0	0	0
2	0	-0.24	0	0	0	0	-0.20
3	-0.30	-0.31	0	-0.20	0	0	-0.28
4	0	0	0	0	0	-0.22	0
5	0	0	0	0	0.27	0	0
6	0	0	0	0	0	0	0
7	0	0	0.29	0	0.20	0	0
8	0.21	0.24	0	0	0	0	0.25
9	0.28	0.32	0	0	-0.26	0	0.35
10	0	0	0	0	0	0	0

Spearman rank correlation coefficient of environment variables with depression for 10 participants over January 2008 to December 2009. Only coefficients with $p < 0.05$ are shown.

A Mood Model

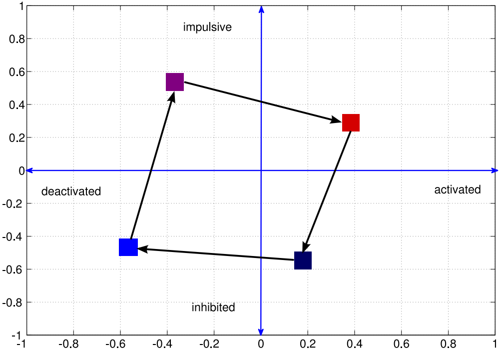

Activation is measured along the x-axis and impulsivity is measured along the y-axis. The rectangles are representative points corresponding to mood episodes, and the path represents a mood trajectory, although not psychologically realistic. We model each mood episode with a hidden Markov model which generates observation vectors on the two rating scales. Transition and observation matrices for the models can be inferred from the data using standard algorithms.

Proposals

Proposals for future work

- Evaluation of exponential smoothing α parameter as a clinical indicator
- Clustering by mood dynamics (Ramoni, *Bayesian Clustering by Dynamics*)
- Controlled dosing for bipolar disorder

Summary

Summary

- Forecasting results even if they are not clinically useful can give an insight into mood dynamics.
- Environmental variables correlate with depression for some patients. Some patients mood correlates negatively with temperature, others positively.
- The data set is unique, and presents many opportunities for further study.