I attended the April 2017 Seismological Society of America (SSA) Annual Meeting in Denver, Colorado, and chose to use the Machine Learning and its Application to Earthquake and Explosion Signal Analysis session as an opportunity to do my seminar summary and analysis. This session was comprised of a series of talks applying machine learning algorithms to problems in earthquake seismology, with topics including from event detection, earthquake clustering, and the application Bayesian autoregressive filters to quantify uncertainties in seismic phase arrivals. But the most interesting and physically compelling talk I heard was Paul Johnson's lecture entitled "Applying Machine Learning to Predict Failure". In this study, *Johnson et al.* used machine learning techniques to study precursors to slip events in a laboratory earthquake machine.

The data set was obtained from a biaxial earthquake simulation apparatus at Pennsylvania State University. This machine imposes a constant normal stress and loading rate to simulated fault gouge (typically glass beads), which exhibits quasiperiodic slip events under the applied load. The machine uses accelerometers to measure time series of dynamic strain in the fault gouge, as well as the applied shear stress, where the latter can be used to time the slip events ("labquakes") during which shear stress rapidly decreases (Figure 1). The objective of the study was to use features extracted from the current dynamic strain time series to predict the time-to-failure of the next slip event. With this predictive model in hand, one could then analyze which subset of features were most predictive, gaining insight into the underlying physical mechanisms driving failure.

To do this, Johnson et al. used a supervised learning technique called a Random Forest (RF). The input data included various functionals of different windows of the strain time series, including the mean, variance, and higher order moments (Figure 1b). The objective of the RF model was to learn the time-to-failure of the following slip event, given these features (Figure 1c). The RF model is comprised of a set of individual decision trees that each use a random subset of the input features to predict a failure time [Murphy, 2012]. While single decision trees are susceptible to overfitting the training data, an ensemble forecast of these trees, each using a randomized subset of the complete set of features, is much more robust. Further, the RF model can be used to isolate which individual features are most diagnostic of the physical failure state of the system.

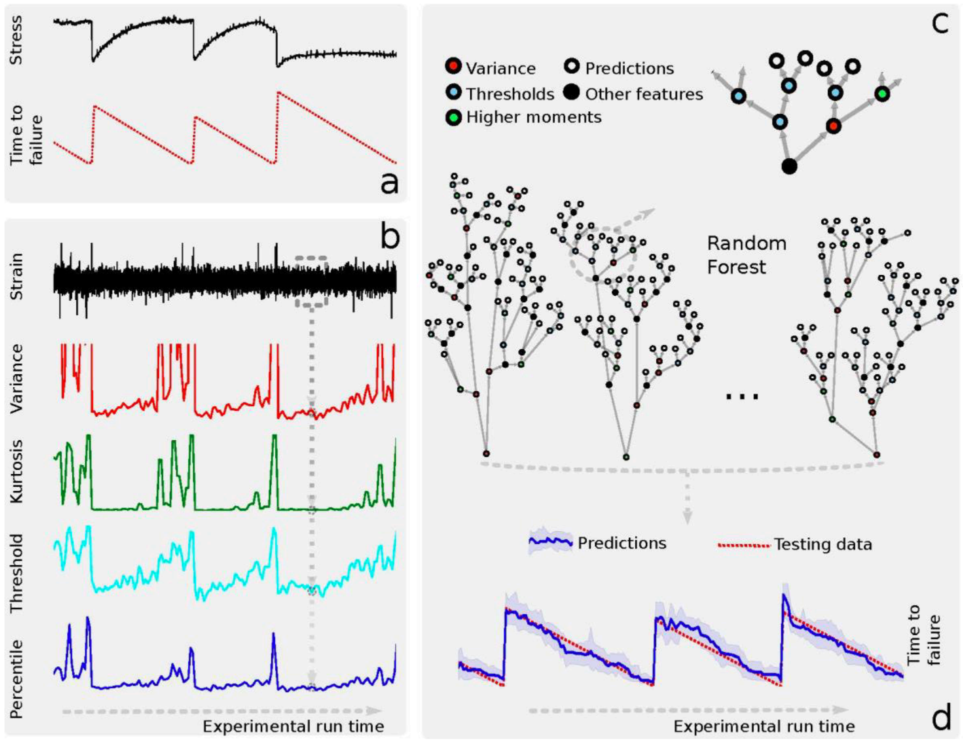

Figure 1: Random Forest (RF) model for predicting time to failure. (a) Experimentally measured shear stress (black), which drops sharply during failure (red). (b) Time series of dynamic strain (black) from the simulated fault gouge, measured by accelerometer. Features input into the RF are derived from windows of this time series. (c) Schematic representation of the RF forest model. (d) RF-predicted time to failure (blue) vs. actual time to failure (red). Figure from *Rouet-Leduc et al., 2017*.

The RF model of *Johnson et al.* proved to be quite effective in predicting failure times (Figure 1d), with the model prediction accounting for more than 90% of the observed variance in failure times ($R^2 \sim 0.89$). Perhaps even more remarkably, one could obtain a reliable prediction of the next failure time shortly after the previous failure, effectively a "long-range" precursor to slip (Figure 2). The most predictive feature turned out to be the variance in the strain time series, where low-amplitude strain signals that resembled tectonic tremor events observed in real fault systems accounted for this variance. This same framework was also applied to study "slow slip" events produced using a different type of gouge material, with similar results. While it is unclear at this point how effective this basic technique would be in forecasting real earthquakes, especially since (i) the lab apparatus has much less aleatory variability than the real earth, and (ii) it is impossible to monitor detailed strain precursors at seismogenic depths (> 5km), the study of *Johnson et al.* is a beautiful example of the power that carefully implemented machine learning techniques can have in helping to characterize and to better understand natural systems and the physical mechanisms that drive them.

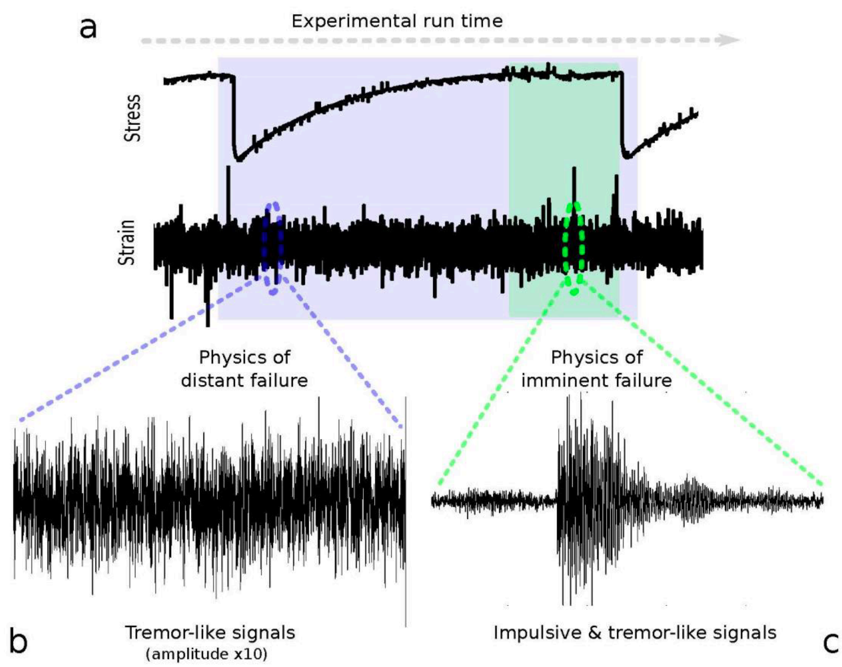

Figure 2: Measured shear stress and dynamic strain preceding a failure event. When failure is distant, the gouge exhibits a persistent but low-amplitude tremor-like signal (b). As failure approaches, tremor becomes impulsive and stronger in amplitude (c). Figure from *Rouet-Leduc et al., 2017*.

References:

1. Johnson, P. (2017). Applying Machine Learning to Predict Failure, Seismological Society of America Annual Meeting (Oral Presentation).

2. Murphy, K (2012). *Machine Learning: A Probabilistic Perspective*, MIT Press, Boston MA, 1104pp.

3. Rouet-Leduc, B. Hulbert, C., Lubbers, N., Barros, K., Humphreys, C. and Johnson, P. A (2017). Machine Learning Predicts Laboratory Earthquakes, arXiv:1702.05774