Fault Detection in Embedded Networked Sensing

Modeling and Computations

Sheela Nair

Department of Statistics
Center for Embedded Networked Sensing
University of California, Los Angeles
snair@stat.ucla.edu

Joint work with Mark Hansen

May 23, 2008

Outline

1. Introduction to Embedded Networked Sensing (ENS)
2. Fault Detection
3. Spatial and Temporal Modeling
4. Signatures and Future Work

Outline

1. Introduction to Embedded Networked Sensing (ENS)
2. Fault Detection
3. Spatial and Temporal Modeling
4. Signatures and Future Work

Embedded Networked Sensing (ENS)

- Spatially distributed network of sensors embedded in the physical world
- Enables spatially and temporally dense environmental monitoring
- Many applications
  1. Terrestrial and aquatic ecosystems monitoring
  2. Seismic and structural monitoring
  3. Contaminant management
  4. Participatory urban sensing

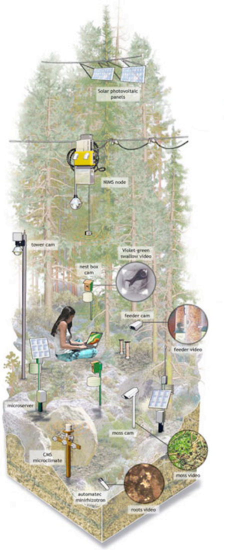

Frank Ippolito,
The New York Times

AMARSS: An ENS Application

- AMARSS deployment set up to examine spatial and temporal distributions of soil characteristics
- 10 stations in an 80 m transect underneath the forest canopy
- Each station: an array of
  1. Above-ground sensors:
     air temperature, relative humidity, and photosynthetic active radiation
  2. Below-ground sensors:
     CO_2 concentration, temperature, water content at 2cm, 8cm, and 16cm.

One Station

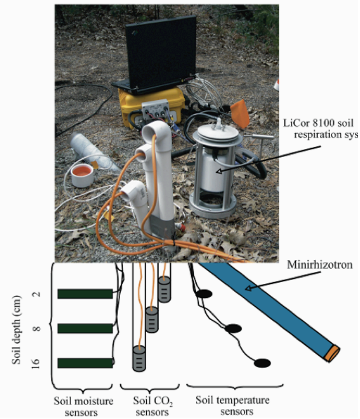

Temperature and CO_2 Data

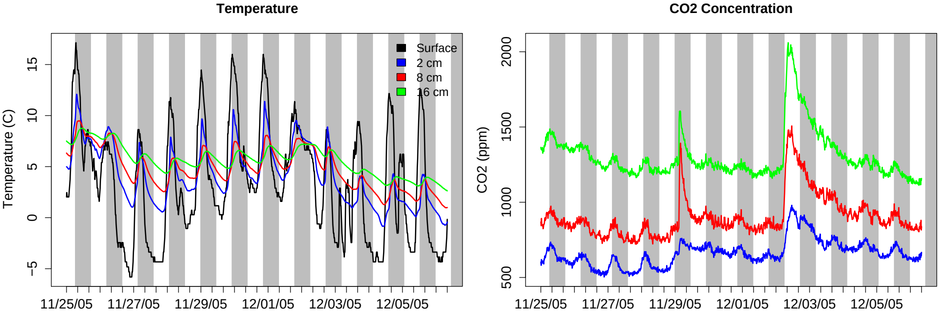

Other Data

- AMARSS itself is part of larger investment in sensing equipment
- James Reserve is equipped with hundreds of microclimate sensors, including weather towers with humidity, ambient temperature, and rainfall sensors
- Data available at sensorbase.org

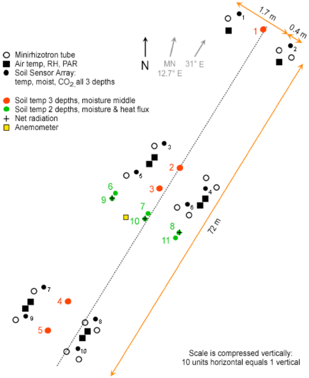

New Themes in Sensor Networks

AMARSS is a representative example of a recent shift in concepts for embedded sensing.

Early Themes

- Visions of thousands of homogeneous, low-powered, autonomous nodes
- Single-purpose, single-user deployments
- Focus on the trade-off between energy and communication constraints

Newer Themes

- Heterogeneous, possibly tiered networks with a rich ecology of mobile and static sensing capabilities
- Multiple experiments and users
- Second-generation questions directly engage data and models to help guide design, deployment, maintenance, and monitoring

Data quality is being recognized as one of the most important challenges.

Outline

1. Introduction to Embedded Networked Sensing (ENS)
2. Fault Detection
3. Spatial and Temporal Modeling
4. Signatures and Future Work

Sensor Faults: Some Examples

- Intermittent outliers
- Stuck-at fault
- Clipping
- Calibration errors
- Sensor drift
- Excessive noise
- Temporary or Total Failure

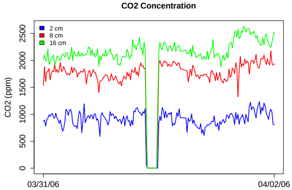

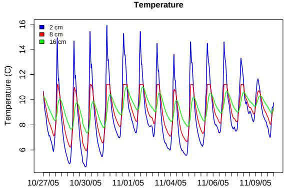

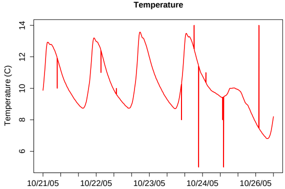

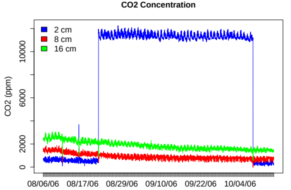

Sensor Faults: Some Examples

- Intermittent outliers
- Stuck-at fault
- Clipping
- Calibration errors
- Sensor drift
- Excessive noise
- Temporary or Total Failure

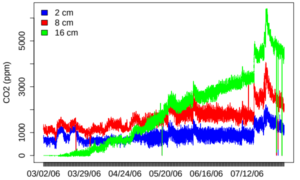
CO2 Concentration

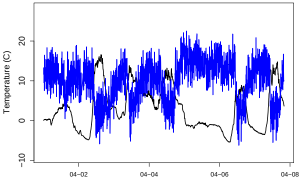
Temperature

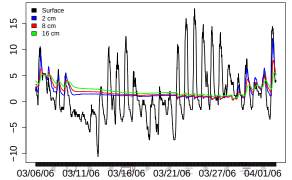
Temperature

Faults and Sources

- Unpredictable environmental conditions
    - Extreme temperatures and precipitation
    - Corrosion
    - Sensor saturation
    - Sensor displacement
- Failures in node hardware and software
    - Unstable sensor connections
    - Low battery power
    - Bugs in software
- Communication errors
    - Noise
    - Compression loss
    - Congestion
- Calibration issues
    - In-field calibration without ground truth
    - Calibration drift over time

Ni et. al (2008) gives a survey of common fault types and sources

Network Topology

- Suppose node s is part of a network
- At time t, node s reports one or more readings:
  $\mathbf{X}(s, t_s) = (X_1(s, t_s), \ldots, X_p(s, t_s))$
- Different nodes can report readings at different times
- Information available to each node (about other nodes) depends on network topology
- This dictates the types of fault detection techniques that can be used

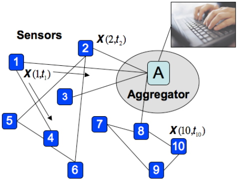

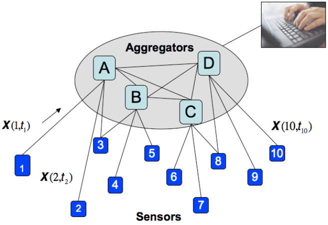

Techniques for Anomaly Detection: A Brief Review

1. Statistical process control from manufacturing industries
2. Fraud detection in telecommunications and banking industries

Statistical Process Control

- Originates with Shewhart's work in the 1920s
- Monitor a manufacturing process using control charts to detect changes
- Shewhart charts, CUSUM charts, EWMA charts, etc.
- Techniques assume that the "in-control" process is iid
- Extensions to temporally correlated data: fit a model to the data and monitor the residuals (we will come back to this later)

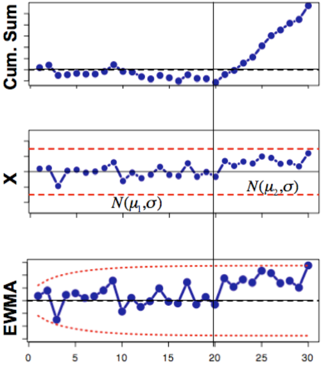

Fraud Detection in Telecommunications and Banking

Considerations:

- Decisions must be made in (near) real-time
- Expert knowledge to help characterize fraudulent behavior
- Many systems based on known rules and thresholding, but the behavior of fraud can change over time

Supervised vs. unsupervised

- Unsupervised: Account behavior is summarized and then compared to itself over time; goal is to detect any deviations from expected behavior (similar to SPC techniques)
- Supervised: Classify account or transaction as legitimate or fraudulent; usually have hard thresholds or decision boundaries
- Signatures: Supervised technique, decisions are account specific and can adapt to changing behavior over time

Outline

1. Introduction to Embedded Networked Sensing (ENS)
2. Fault Detection
3. Spatial and Temporal Modeling
4. Signatures and Future Work

Model-Based Approach to Anomaly Detection

Basic Idea

1. Develop models using clean training data
2. Use training data to get prediction error and interval
3. At time t, use model to predict $\hat{Y}(t+1)$
4. If $(Y(t+1) - \hat{Y}(t+1)) \notin$ Prediction Interval
    - flag the observation
    - substibute $Y(t+1) = \hat{Y}(t+1)$

Model-Based Approach to Anomaly Detection

Basic Idea

1. Develop models using clean training data
2. Use training data to get prediction error and interval
3. At time t, use model to predict $\hat{Y}(t+1)$
4. If $(Y(t+1) - \hat{Y}(t+1)) \notin$ Prediction Interval
    - flag the observation
    - substibute $Y(t+1) = \hat{Y}(t+1)$

Considerations

- Exploit physical model whenever possible
- Keep in mind where the computation will "live" and what information is available at each level in the network
- Cannot store historical data, must update the models online

Time Series Models: Related Work

Fitting univariate time series models to sensor network data on a single node has received quite a bit of attention in the literature, mostly in the context of energy savings.

Related work:
Tulone and Madden (2006) fit local AR(p) models fit at each node. Models are refit when predictions are systematically too far from the truth.

Le Borgne et. al (2007) do adaptive model selection to choose order of the models by minimizing the amount of communication between the node and sink.

Online AR models work well for the soil temperature data.

Temporal Model: Back to Soil Temperature Data

Suppose we want to detect faults at the single sensor level.

Let:

$$\begin{aligned} Y(t,d) &\equiv \text{temperature at time } t \text{ on day } d \\ \mu(t) &\equiv \text{a mean function that varies with time of day.} \end{aligned}$$

AR(p) model for Y

$$[Y(t+1,d) - \mu(t+1)] = \sum_{j=0}^{p-1} \phi_j [Y(t-j,d) - \mu(t-j)] + \epsilon(t+1,d)$$

- The AR model can vary with time
- Many ways to use historical data to estimate $\mu(t)$
- Computational considerations are important

Temporal Model: Detecting Faults

Outlier Faults

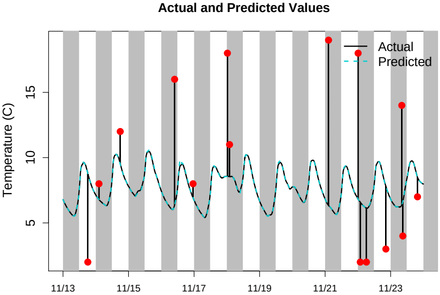

Stuck-At Faults

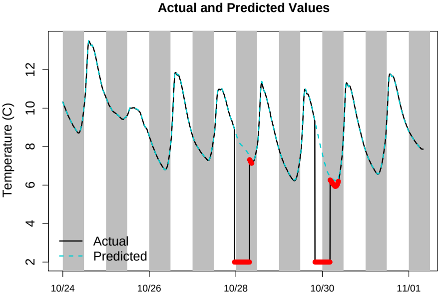

Time series model is able to detect faults that cause large gradients.

Temporal Model: Detecting Faults

Sensor Drift

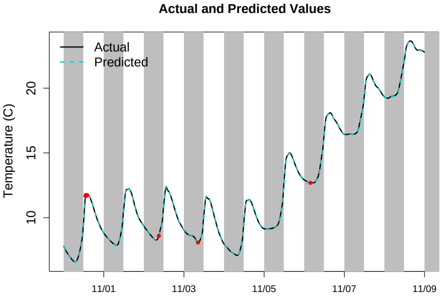

The model learns the slow drift and cannot detect the fault.

Data from Multiple Sensors

So far, we have been dealing with data from a single sensor over time.

Depending on the network topology, we can develop “spatial” models for data from multiple nodes.

We will borrow from some physical models describing how temperature diffuse through soil.

Spatial Model for Soil Temperature Data

Physical model for annual soil temperature at depth s (Hillel 1982):

$$Y_s(d) = T_a + A_0 e^{-s/c} \sin\left(\frac{2\pi(d-d_0)}{365} - \frac{s}{c} - \frac{\pi}{2}\right)$$

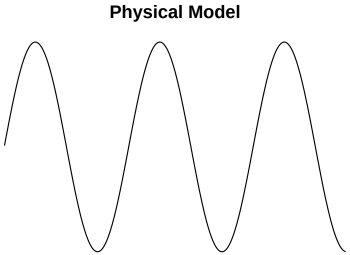

Spatial Model for Soil Temperature Data

Physical model for annual soil temperature at depth s (Hillel 1982):

$$Y_s(d) = T_a + A_0\, e^{-s/c} \sin\left(\frac{2\pi(d-d_0)}{365} - \frac{s}{c} - \frac{\pi}{2}\right)$$

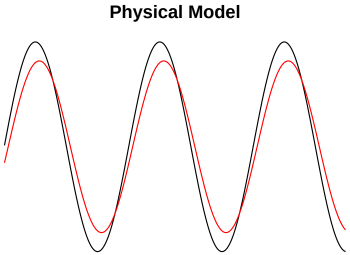

Exponential damping and lag that depends on depth s and c, which depends on thermal diffusivity of the soil

Spatial Model for Soil Temperature Data

Physical model for annual soil temperature at depth s (Hillel 1982):

$$Y_s(d) = T_a + A_0 \, e^{-s/c} \sin\left(\frac{2\pi(d - d_0)}{365} - \frac{s}{c} - \frac{\pi}{2}\right)$$

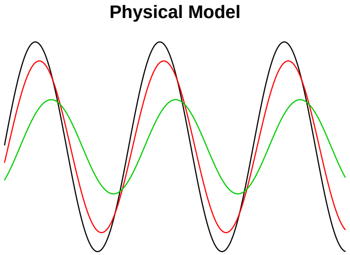

Exponential damping and lag that depends on depth s and c, which depends on thermal diffusivity of the soil

Spatial Model for Soil Temperature Data

Physical model for annual soil temperature at depth s (Hillel 1982):

$$Y_s(d) = T_a + A_0\, e^{-s/c} \sin\left(\frac{2\pi(d - d_0)}{365} - \frac{s}{c} - \frac{\pi}{2}\right)$$

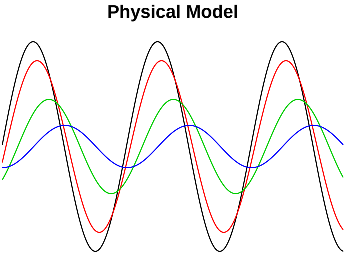

Exponential damping and lag that depends on depth s and c, which depends on thermal diffusivity of the soil

Spatial Model for Soil Temperature Data

Physical model for soil temperature at depth s (Hillel 1982):

$$Y_s(d) = T_a + A_0 \, e^{-s/c} \sin\left(\frac{2\pi(d - d_0)}{365} - \frac{s}{c} - \frac{\pi}{2} \right)$$

How does this relate to our higher-resolution data?

Spatial model based on physical model:

$$Y_s(t) = \beta_{0,s}(t) + \beta_{1,s}(t) \, Y_0(t - k_s(t)) + \epsilon_s(t)$$

Spatial Model for Soil Temperature Data

Spatial model:

$$Y_s(t) = \beta_{0,s}(t) + \beta_{1,s}(t) Y_0(t - k_s(t)) + \epsilon_s(t)$$

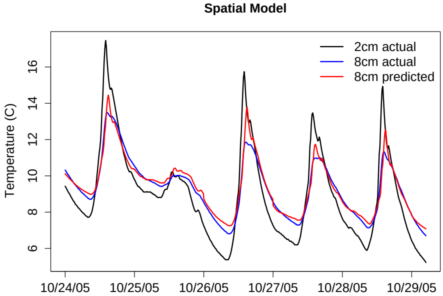

In this simple case, the model does not completely capture the smoothing that happens with depth. This is the subject of current work.

Spatial Model: Detecting Faults

Sensor Drift

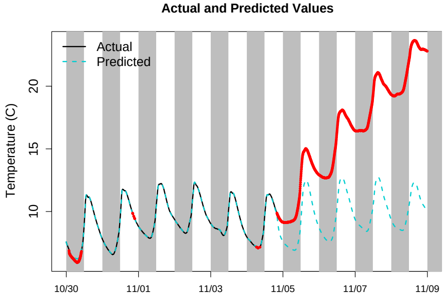

Given information from properly functioning sensor located at higher depth, the model can detect drift.

Clipping Faults

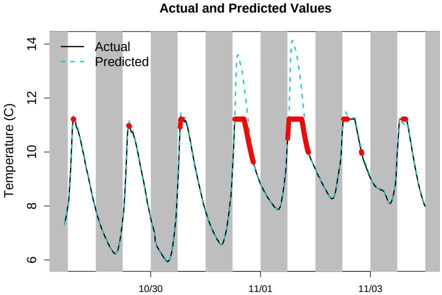

Catches some of the clipping faults, but then substitution of $\hat{Y}(t)$ results in many false positive flags.

A Bigger Picture

Work so far is data-driven. The picture is much bigger than this...

- Different users of the deployment have different needs and priorities
- We usually have subject matter knowledge about both the application and the types of faults common sensor faults
- Difficult to leverage expert knowledge into purely model-based techniques
- Signatures provide a middle ground between rule- and model-based techniques

Outline

1. Introduction to Embedded Networked Sensing (ENS)
2. Fault Detection
3. Spatial and Temporal Modeling
4. Signatures and Future Work

Building and Maintaining Sensor Signatures

Signatures for monitoring multidimensional behavior over time. Algorithm adapted from Cahill et. al (2000).

- F_n is multivariate distribution of M features after n observations have been received
- Can write F_n as a product of the conditional distributions: $F_n(\mathbf{X_n}) = F_n(X_{1,n})F_n(X_{2,n}|X_{1,n}) \ldots F_n(X_{M,n}|X_{1,n} \ldots X_{M-1,n})$
- Signature is fixed-length estimate of F_n.
- Express signature as a product of conditional distributions: signature components
- Signature components represented by univariate density estimates.

Building and Maintaining Sensor Signatures

Let f_0 be a sensor signature, f_1 be a fault signature, and $\mathbf{X}$ be M features that distinguish between normal and faulty behavior.

Procedure to Detect Faults

When the $(n+1)$th observation $\mathbf{X}_{n+1}$ is received:

1. Calculate a score: $S(n+1) = \log \frac{f_{1,n}(\mathbf{X}_{n+1})}{f_{0,n}(\mathbf{X}_{n+1})} = \sum_{i=1}^{M} \log \left(\frac{f_{1,n}(X_{i,n+1})}{f_{0,n}(X_{i,n+1})} \right)$
2. If $S(n+1) \leq 0$:
   observation is not faulty, update $f_{0,n+1}$ using $f_{0,n}$ and $\mathbf{X}_{n+1}$

   else:
   flag $X(n+1)$ probabilistically

- Makes use of known fault patterns and other expert knowledge about the application
- Tests specifically for anomalies that look like faults
- Allows for sensor-specific thresholding since each observation is compared to its own sensor's signature

Ongoing Work Using Signatures to Detect Sensor Faults

- Collaborating with biologists to identify set of features
- Incorporating work with spatial and temporal models in sensors signatures
- Exploring whether we can gain additional information about fault type from scores of flagged observations

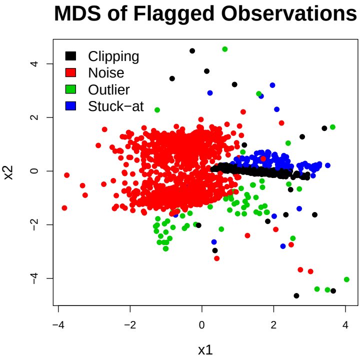

Conclusions

- Assessing data quality is a crucial part of the process of embedded sensing.
- Deployments have a lot of information we can draw on in a fault detection system. We must keep the network topology in mind when designing systems for ENS.
- Deployments are multi-user. We not only want to detect faults but also try and determine the case and severity of the problem.
- Signatures are somewhere in the middle between model- and rule-based detection schemes. They provide a nice way to incorporate subject matter knowledge.

Sources

- Cahill, M. H., Lambert, D., Pinheiro, J. C., and Sun, D. X. (2002), "Detecting Fraud in the Real World," *The Handbook of Massive Data Sets*, Kluwer Academic Publishers, 911 - 929.
- Le Bornge, Y., Santini, S., and Bontempi, G. (2007), "Adaptive Model Selection for Time Series Prediction in Wireless Sensor Networks," *Signal Processing*, 87(12), 3010 - 3020.
- Hillel, D. (1982), "Introduction to Soil Physics," Academic Press, San Diego, CA.
- Ni, K., Ramanathan, N., Hajj Chehade, N. M., Balzano, L., Nair, S., Zahedi, S., Pottie, G., Hansen, M., and Srivastava, M. (2008), "Sensor Network Data Fault Types," *ACM Transactions on Sensor Networks* (to appear).
- Tulone, D. and Madden S. (2006), "An Energy-Efficient Querying Framework in Sensor Networks for Detecting Node Similarities", *In Proceedings of the 9th International ACM Symposium on Modeling, Analysis, and Simulation of Wireless and Mobile Systems.*