Numenta

Geospatial Tracking

Learning the Patterns in Movement and Detecting Anomalies

White Paper

Executive Summary

This white paper describes an application, developed by Numenta, for geospatial tracking and anomaly detection. This application utilizes Hierarchical Temporal Memory (HTM), a biologically inspired computational theory based on the neocortex, to automatically model typical travel patterns and identifies anomalies in movement. Key points covered in this white paper are as follows:

- This application detects the patterns and anomalies in movement and speed.
- Applications spanning personnel monitoring, fleet management, logistics, smuggling prevention and air safety stand to benefit from the enhanced geospatial tracking capabilities offered.
- Automatic modeling of travel patterns at the unit-level provides a heightened level of flexibility, scalability and precision compared to existing geospatial monitoring and alerting techniques.
- Travel patterns are learned and anomalies are detected by encoding GPS data plus speed into Sparse Distributed Representations (SDRs) that are then modeled by the Cortical Learning Algorithm (CLA).
- Examples are provided to show anomalies that have been detected based location, speed and route pattern:
  - Location – unexpected path deviations, new routes taken, detours.
  - Speed – irregular slowdowns, bottlenecks, unexpected travel modes (driving speed versus walking).
  - Route Patterns – abnormal sequences of events within previously learned geographical areas such as U-turns, repeat stops, etc.
- Parties interested in recreating or expanding upon these results can access the encoder and algorithms used in this application through NuPIC, Numenta's open source project.
- This application takes advantage of the same underlying Hierarchical Temporal Memory theory and Cortical Learning Algorithm code base that is used in other Numenta applications, demonstrating the generalizability of Numenta's technology to a variety of use cases.

Advances in Anomaly Detection

a·nom·a·ly	Something that deviates from what is standard, normal, or expected.

The ability to detect anomalies in real time can be incredibly valuable. Imagine being able to notice early warning signs for failure in a large turbine, see slight variations in a heartbeat indicating disease, or detect an unusual pattern of shopping cart failures on an ecommerce website. Anomalies are not always bad or indicative of a failure. For example, detecting a subtle change in consumer buying habits could provide an opportunity to discover a new trend. In today's world where the amount of data being collected is exploding, the opportunity for detecting anomalies is rapidly growing.

However, accurately detecting anomalies can be very difficult. First, what qualifies as an anomaly is always changing. Systems evolve over time as software is updated or as behaviors change. Therefore, effective anomaly detection requires a system to learn continuously. Second, to detect anomalies early one can't wait for a metric to be obviously out of bounds. Early detection requires the ability to detect subtle changes in patterns that are not obvious or easily detected. Furthermore, because anomalies by their nature are unexpected, an effective detection system must be able to determine whether new events are anomalous without relying on preprogrammed thresholds.

At Numenta we have taken a fresh approach to this problem and have created what we believe is the world's most powerful anomaly detection technology. This approach is derived from our understanding of the neocortex, which is itself a powerful prediction and anomaly detection system. Our suite of applications takes advantage of this understanding to drive state-of-the-art breakthroughs in two dimensions: 1) how we utilize the processes of the brain to model data, and 2) how we detect anomalies based on that model. This paper will describe these advances by illustrating how Numenta's geospatial tracking application detects anomalies in location, speed and movement patterns.

Geospatial Tracking and Anomaly Detection

A wide range of applications stand to benefit from an improved ability to analyze movement and travel patterns. Aircraft controllers charged with monitoring the locations and flight paths of hundreds of airplanes can receive a warning at the earliest hint of a plane diverting from its typical trajectory. Parents concerned with the safety of their families can use this application to monitor the movement of their children between home, school, sports and play areas. This application automatically learns normal routes and modes of travel, and sends alerts if children move outside the range of their expected locations and behaviors. Logistics companies can also use this application to track the movement of inventory across their supply chain networks – typical shipping methods or routes can be learned, and operations managers can be notified as soon as delays arise. These and other potential use cases are summarized below:

- Air Safety
  - Monitor location, speed and trajectory of commercial aircraft
- Maritime Domain Awareness
  - Detect abnormal ship routes or unexpected ports of entry
  - Spot unexpected travel impediments – weather, malfunctions, piracy
- Family and Pet Care
  - Receive alerts when family members deviate from normal routes or locations
  - Track teen driving patterns and speeds
- Fleet and Professional Driver Tracking
  - Monitor driver route and activity compliance
  - Collect anomalous route data to understand traffic conditions or road obstructions

- Personnel Compliance and Safety
    - Monitor employee movements in hazardous industries - mining, public safety, military
    - Ensure service coverage - security guards, hospitality workers, sales associates
- Supply Chain Visibility
    - Monitor inventory movement through supply chains to identify anomalies and delays
    - Spot new bottlenecks or service disruptions in delivery networks

Most geospatial tracking and alerting systems in use today rely on unsophisticated geo-fencing methods. Geo-fencing requires users to define specific geographical boundaries (or thresholds) that separate safe areas from areas that trigger alerts. Setting up these boundaries is very time consuming and often imprecise. Additionally, as circumstances on the ground change, effort must be spent to adjust these thresholds so that they don't send out false-positive (or false-negative) alerts. Other techniques used to monitor location and travel in the logistics industry only rely on point-to-point scanning. While this enables inference of movement patterns between points, no data actually exists to identify bottlenecks, disruptions or intermediate points of transfer.

The Numenta Geospatial Tracking application offers significant improvements upon the existing methods of location tracking and alerting across the following dimensions:

Setup	Automatic model building and learning eliminates the need to manually define and maintain boundaries.
Scalability	Unique models can be built automatically for each person or object being monitored.
Precision	Models are developed that monitor specific travel paths and speeds at every route segment – users no longer have to make broad generalizations about expected geographies and circumstances.
Flexibility	Online learning enables continual model updating to account for new routes, alternative paths, or changes in speed.
Scope	Ability to identify travel anomalies based on rate of travel, direction and order of movement, in addition to location.

Modeling, Prediction and Anomaly Detection

The Numenta Geospatial Tracking application excels at rapidly learning patterns in movement and detecting abnormalities in the direction, route and speed traveled. Figure 1 below summarizes the key steps Numenta uses to model geospatial data and identify anomalies.

Figure 1 - Process for Modeling Movements and Identifying Anomalies

Data is collected periodically that includes current GPS location, speed and a time stamp. This data is fed to an encoder that converts each data snapshot into a Sparse Distributed Representation (SDR). SDRs are the language of the brain and enable several useful attributes such as generalizability across data stream type, strong resistance to noise, and the attachment of semantic meaning to data points. More information can be found on Numenta's website or through NuPIC, Numenta's open source project that contains descriptions and implementations of the algorithms used in our applications.

To create an SDR, the encoder maps a grid of bits across the map area in which an object is moving. When the tracked object is moving at a slow speed, the encoder will create an SDR that contains a geographically small snapshot of the area around the object (Figure 2a). Alternatively, the SDR created for an object moving rapidly at the same exact location would contain represent a larger, more disperse geographical area (Figure 2b). In each case, the number of "on" bits will be the same. As a result, the SDR generated for a tracked object moving at a walking speed of 5 mph at a given point will be different than the SDR generated for an object being transported in a car at the exact same location, although there will be many bits of overlap.

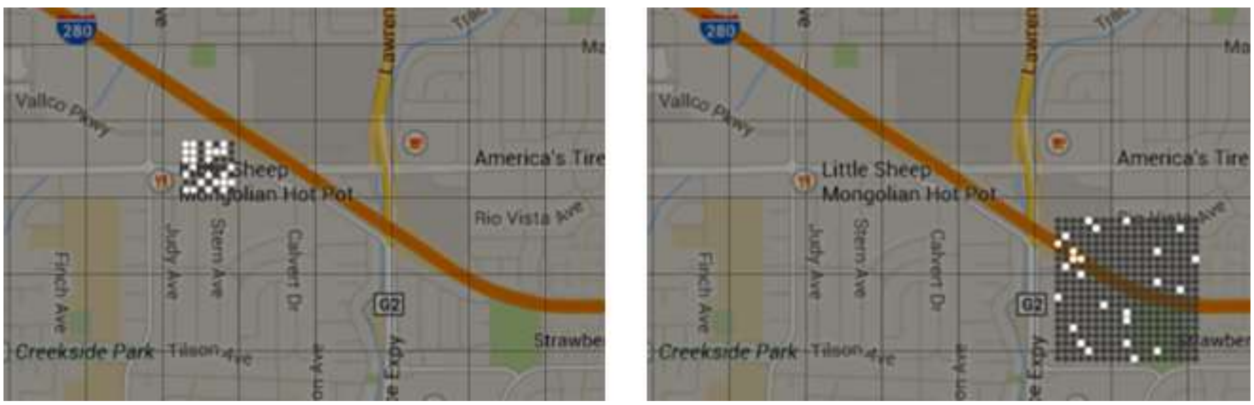

Figure 2 - Encoding of an SDR based on location and speed: (a) Slow movement; (b) Fast movement

Next, a sequence of SDRs is fed into the Cortical Learning Algorithm (CLA). The CLA is a simulation of a small slice of the neocortex, and is responsible for learning the patterns and sequences of a user's movement. Just as the brain is continually thinking ahead about the next step in one's path, the CLA is constantly predicting where a tracked person or object will be once the next data snapshot is taken. When each new data point arrives, the CLA compares its prediction and level of confidence to the new input to see how correct its prediction was. As a result, every new data point results in a scalar estimate of the accuracy of the previous prediction. This value, known as the raw anomaly score, is then scaled according to its relative position in the normal distribution of aggregated anomaly scores recorded over a longer time period.

Learning a Route

Before anomalies can be detected, the Numenta Geospatial Tracking application must first learn what is normal. Directional knowledge of a route along with an initial understanding of the likely speed traveled throughout the trip can be learned after a small number of outings. Subsequent outings reinforce the predictability of the route and tune the application's expectations for speed traveled at each point along the path.

Routes traveled along with anomaly data are superimposed on a map for intuitive visualization. Lines overlaid on the map show the travel path taken, while route markers, or chevrons, indicate the speed and direction of travel. The color of these lines and markers represent the degree to which a location, speed or directional anomaly was detected. A red route overlay or chevron will indicate a highly anomalous route segment. Yellow and green colors represent progressively more predictable paths, speeds and directions. Additionally, since measurements are taken at a fixed interval, the amount of space between route markers denotes the speed traveled – closely packed route markers indicate slow speeds of travel while increasingly more spread out markers indicate faster speeds.

Figure 3 below shows the anomaly scores for a route taken between Cupertino, CA and Redwood City, CA. The image on the left shows the results from the first time this route was monitored. Because the route has not yet been learned, the entire trip shows up as unpredictable or anomalous. The middle image represents the application's second time monitoring this route. This image shows a much higher degree of predictability with significantly fewer red markers. By the third time this route was taken (right image), the trip has been learned. Subsequent trips will continue to tune the application's expectations for typical travel patterns, traffic slowdowns and alternative paths taken along this route.

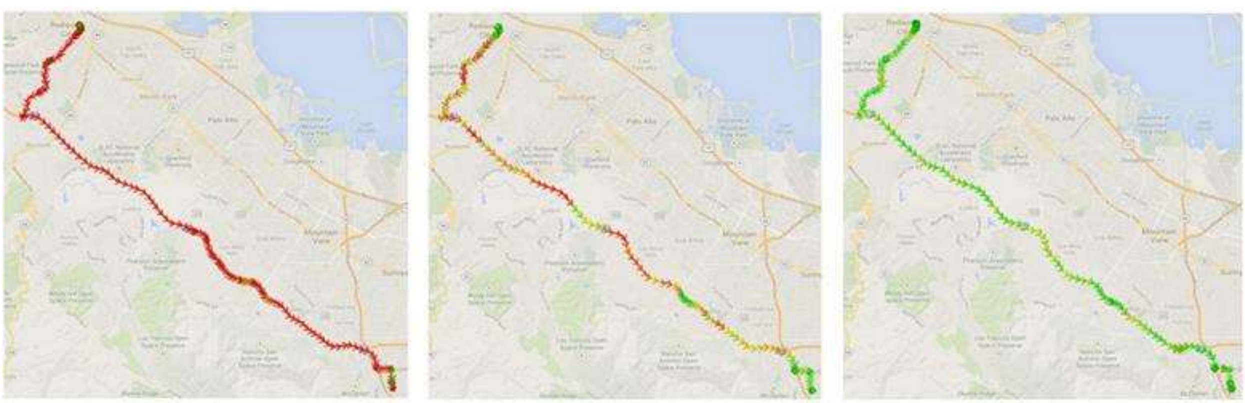

Figure 3 - Learning a New Route: (a) First trip on new route; (b) Second trip; (c) Third Trip, showing traffic anomalies

Geospatial Tracking in Practice

This application makes it easy to identify anomalies in the movements of the people or objects being monitored. For example the dashboard below in Figure 4 represents the data collected during a commute from Cupertino, CA to Redwood City, CA over the course of a week. This image shows a normal route without any significant anomalies. Compare this to the trip shown next in Figure 5 where an anomaly was detected due to the commuter taking an unscheduled stop. By drilling down into this data, one can easily view the details of the irregular portion of the route in order to investigate the abnormal activity.

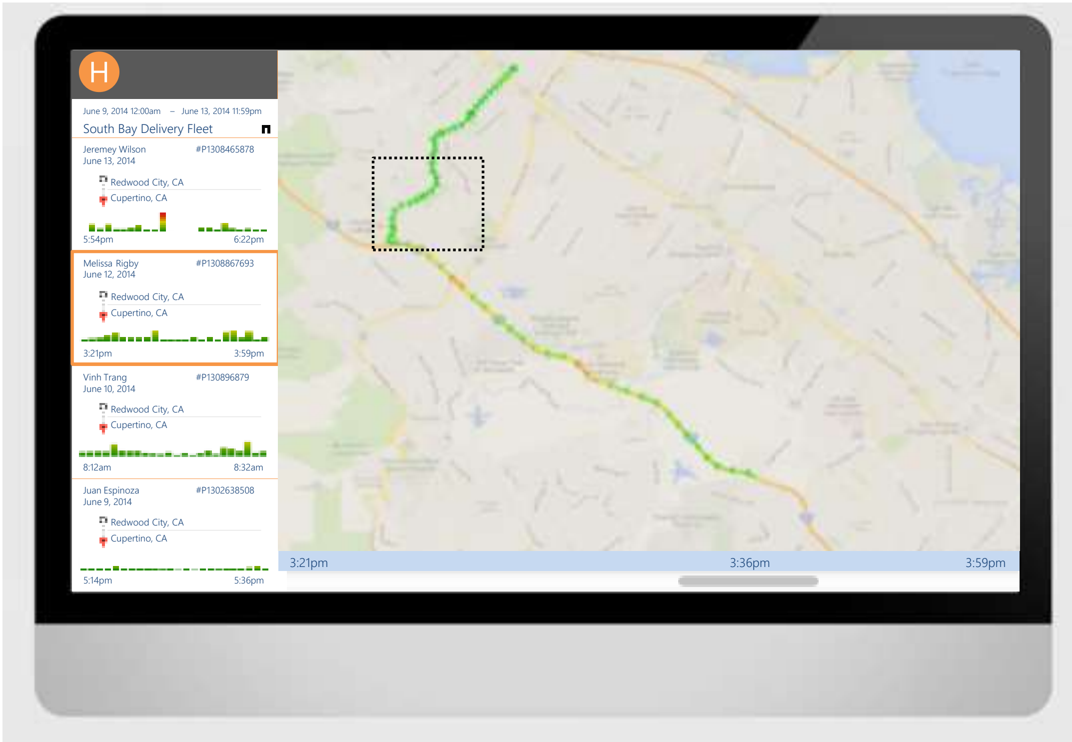

Figure 4 - Normal route with no anomalies in outlined area

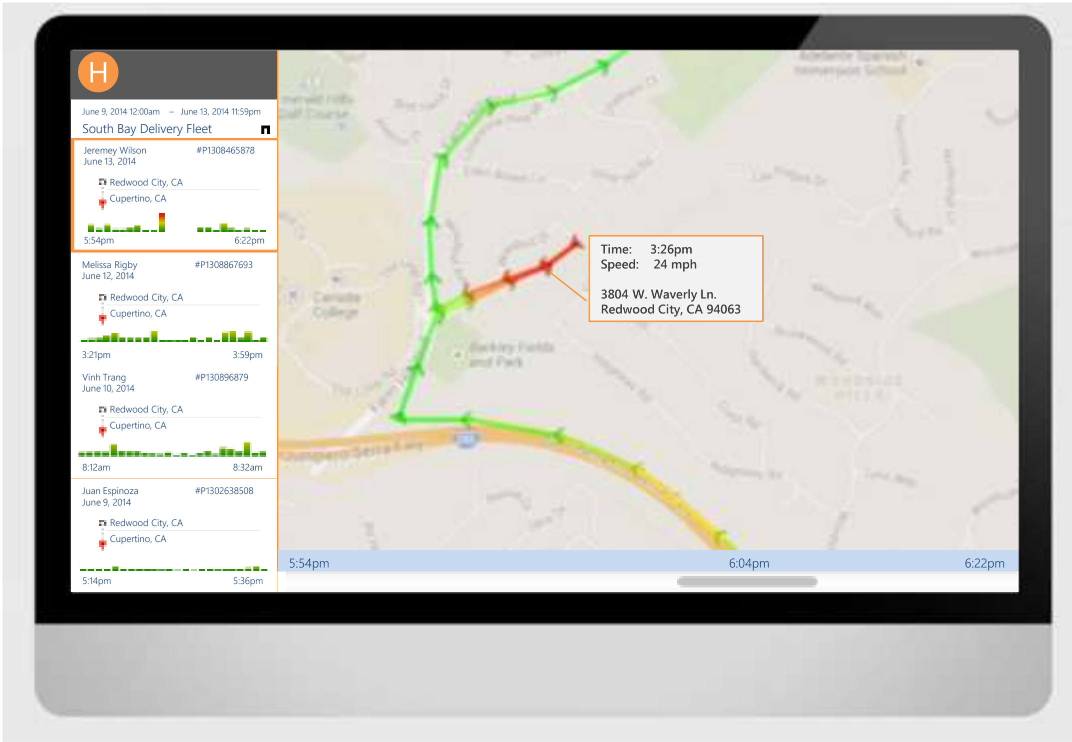

Figure 5 – Portion of the route shown in Fig. 4 showing an anomalous deviation in the typical path

The ability to identify unscheduled stops like these can be very valuable for the managers of delivery fleets or field service operations who are concerned with driver adherence and operational efficiency. Additionally, coast guards or navies may utilize this same capability to identify unexpected calls on port made by ships engaged in the smuggling of illegal contraband.

Another example that highlights this application's ability to identify abnormalities in travel patterns is shown below. Figure 6 and 7 show the two standard paths taken by an individual walking to get coffee. Notice that although these two paths are different, no anomalies are displayed because the application already considers both of these paths to be normal. Contrast this with Figure 8 that shows the same individual beginning to walk along the normal travel path; however, instead of continuing on to the cafe, this individual gets into a vehicle at point A. At this point, the application immediately flags an anomaly due to the abnormal speed of travel, even though the initial portion of this route followed along the exact same path as was typically walked.

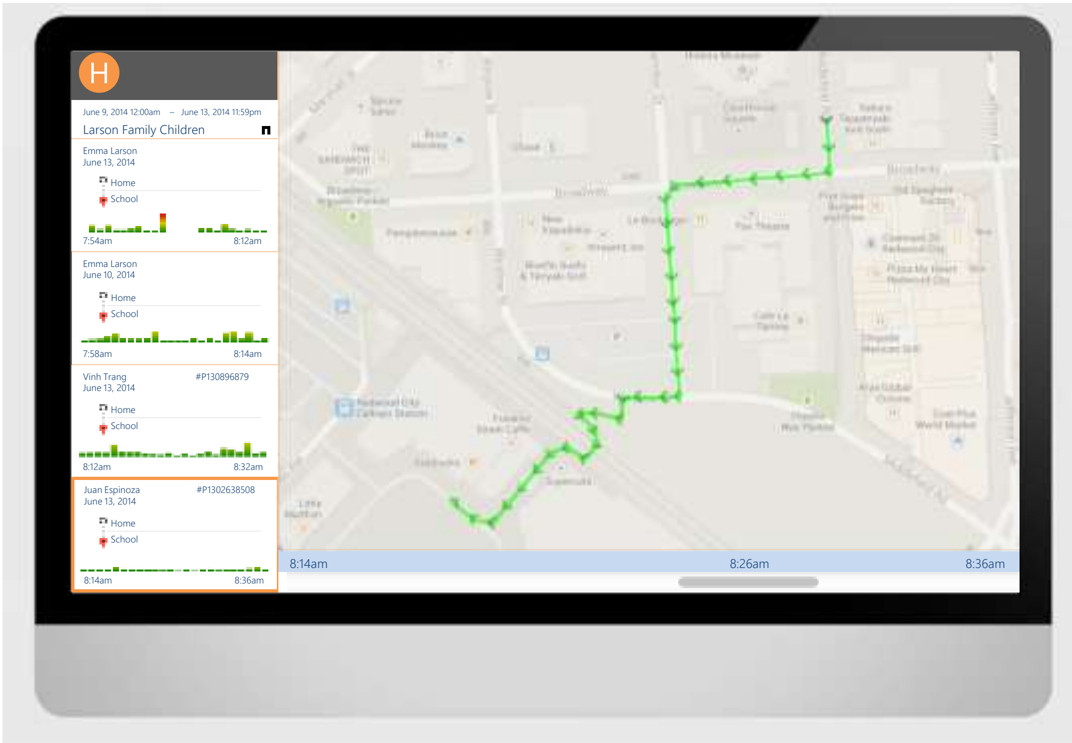

Figure 6 - Standard route individual takes from work to the cafe

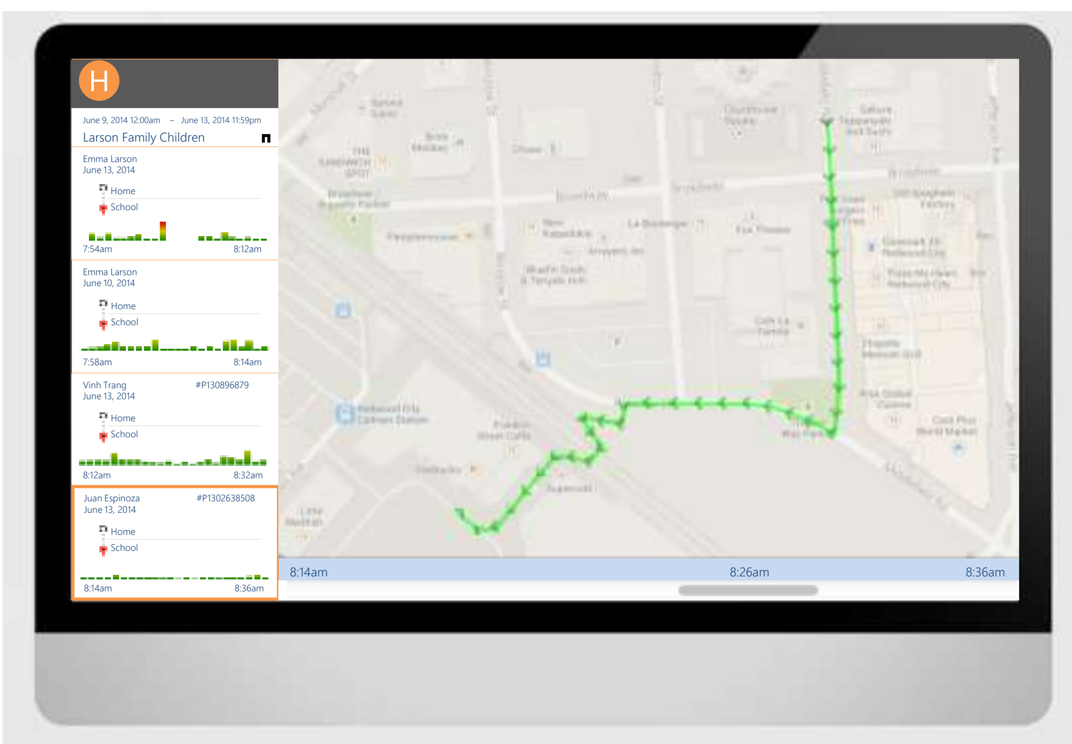

Figure 7 - Alternative, yet previously learned route taken from work to cafe

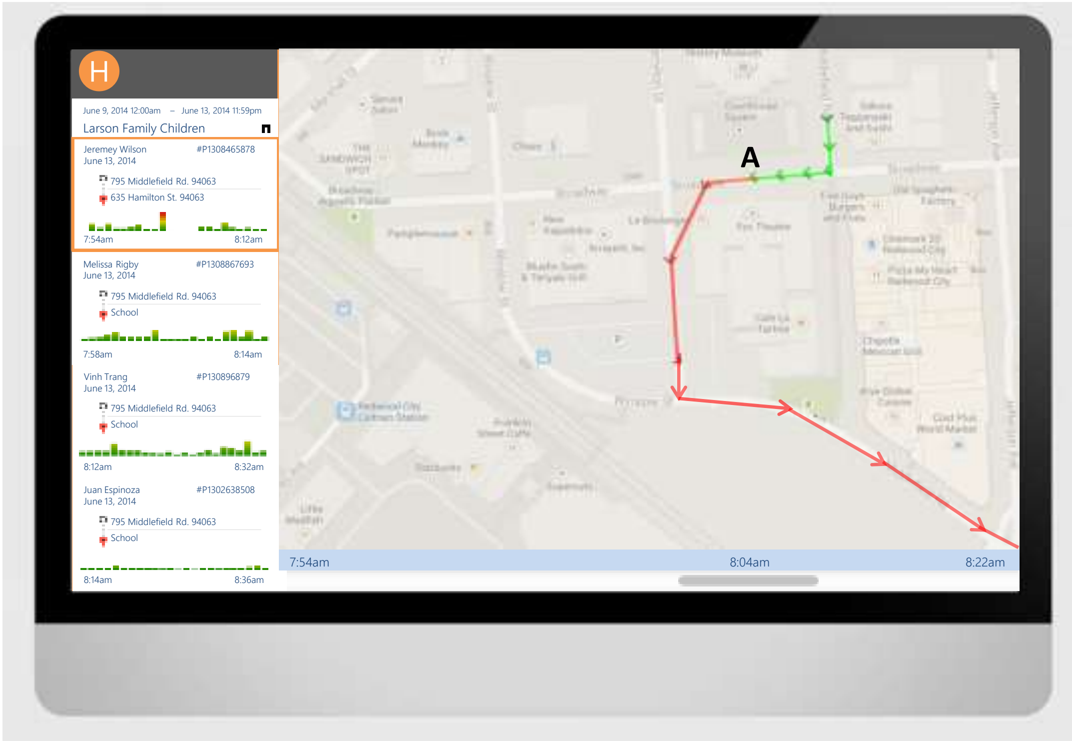

Figure 8 - Anomaly detected immediately at point A due to abnormal speed

This example illustrates the intuitive and flexible manner with which this application can monitor the movements of people, vehicles and other objects. Its ability to automatically learn normal movement routines eliminates the need for extensive system setup and maintenance required by other geo-monitoring technologies. Furthermore, the application's ability to understand alternative paths and differentiate between expected and unexpected travel patterns significantly reduces the generation of erroneous alerts, enabling a very precise level of monitoring.

A final scenario that underscores this application's advanced capabilities is shown below in Figure 9. Here, an individual was monitored while traveling along their normal commute to work. Along the way, the individual suddenly made a U-turn and began travelling backwards along same street. Even though this individual never deviated from their normal path, the application started generating anomalies once the individual started to travel in the opposite direction. As soon as the individual began traveling in the normal direction again, the generation of anomalies stopped.

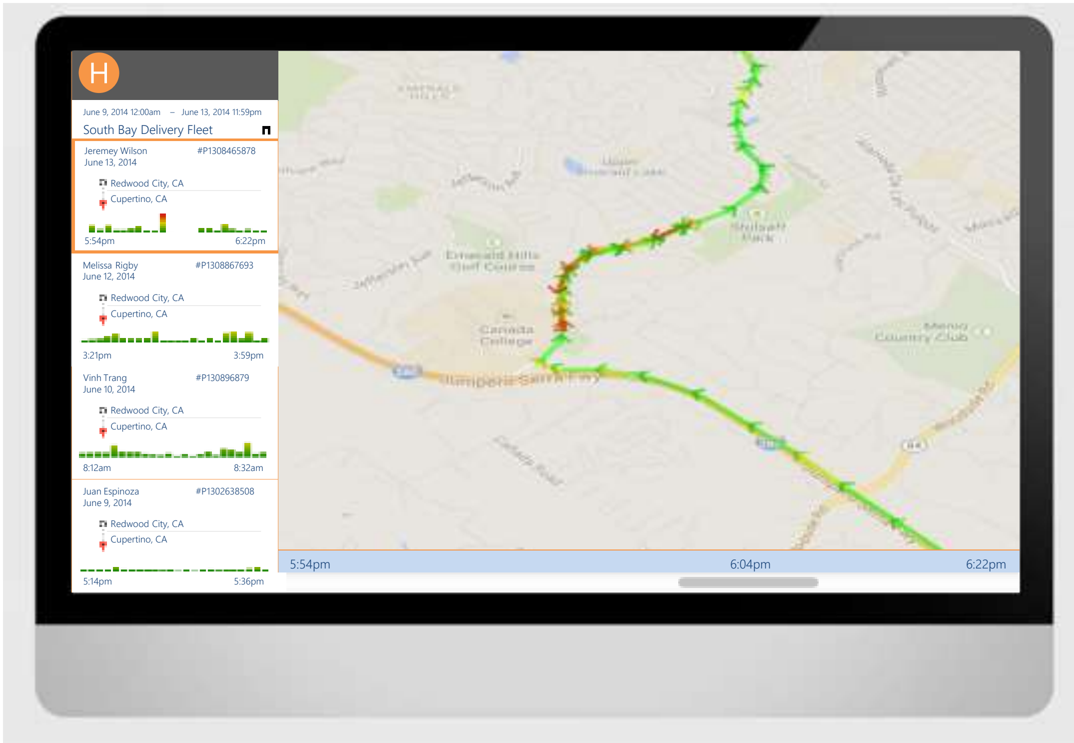

Figure 9 - Wrong direction on path flagged as anomaly

Geo-fencing and other standard location tracking technologies are not capable of identifying anomalies like these because the monitored individual never left the standard path. However, this application's ability to understand the temporal sequences in movement allow it to note anomalies such as these with ease.

Conclusion

The simplicity and precision of the Numenta Geospatial Tracking application provide a host of new opportunities that advance the practice of location monitoring and movement anomaly detection. Underlying these advances is the ability to automatically learn and predict the fundamental patterns of movement, speed and direction that make up more complex travel paths. This gives Numenta's Geospatial Tracking application the ability to detect a much wider range of anomalies than is possible with crude geo-fencing, while simultaneously minimizing the generation of false-positive (or false-negative) alerts. Furthermore, the sophisticated way that this application is able to learn new movement patterns and separate true anomalies from minor deviations gives its users a flexible, scalable and intuitive way to stay up to speed with their geospatial monitoring requirements in a continually evolving world. This gives its users a relevant, real time understanding of the movement patterns of the people and objects they care about.

Although this application makes the practice of geospatial tracking and movement monitoring simple and easy to use, the science behind this technology is profound. This application takes advantage of the same underlying Hierarchical Temporal Memory theory and Cortical Learning Algorithm code base that also underpins Numenta's other applications (Figure 10). If you are interested in recreating the results shown above or developing new applications that take advantage of these advances in geospatial tracking, head to NuPIC to download this application's source code and geospatial encoder.

To learn more, we invite you to read about the neuroscience and computer science used on our website. By applying years of research in neuroscience and computer science, we believe that our approach to anomaly detection and pattern recognition represents a significant step forward for the monitoring of anything that generates continuous data.

Figure 10 - The same HTM code base underpins Numenta's diverse suite of applications

About Numenta

Numenta was founded in 2005 to lead a new era of machine intelligence. Numenta builds solutions that help companies automatically and intelligently act on data. Its biologically inspired machine learning technology is based on a theory of the neocortex first described in co-founder Jeff Hawkins' book, *On Intelligence*. This technology is ideal for the analysis of continuously streaming data sets and excels at modeling and predicting patterns in all types of data. The application described above is just one of a suite of products and applications the utilize Numenta's flexible and generalizable Hierarchical Temporal Memory machine learning technology to provide solutions that encompass the fields of machine generated data, human behavioral modeling, geo-location processing, semantic understanding and sensory-motor control. In addition, Numenta has created NuPIC (Numenta Platform for Intelligent Computing) as an open source project. Numenta is based in Redwood City, California.