Trace

Prototype Phase Proposal

Nittany AI Competition 2021

Project Members

Zi Xian Leong — University Park — College of Earth and Mineral Sciences

Junzhu Shen — University Park — College of Earth and Mineral Sciences

Ian Lee — University Park — College of Earth and Mineral Sciences

Robert Krimetz — University Park — Eberly College of Science

Raymond Friend — University Park — Eberly College of Science

Prototype Video

https://www.youtube.com/watch?v=moRX4Hfn-Yo

Primary Contact Name:
Zi Xian Leong

Email:
zul311@psu.edu

Phone Number:
(206) 669-4633

Supervisor Name:
Brad Zdenek

1 Problem Statement and Project Overview

According to Our World in Data, natural disasters kill on average 60,000 people per year, globally; and disasters disproportionately affect the poorest of people [1]. And natural disasters do not spare large cities. Cities face a double problem related to natural disasters: a large number of people for which to care, and too little space and too few resources to always evacuate everyone during a natural disaster. An infamous example of a city's failure to care for its population upon the advent of a natural disaster is New Orleans, LA during Hurricane Katrina in 2005 [2]. Residents remained stranded in their attics, on their roofs, or on the streets for days before being rescued because of the flood, and not everyone had the same opportunity to escape the city before its destruction. Such shortcomings could be called *un*-natural disasters: inadequate response to human needs. Cities have too few urban monitoring tools to prepare for, handle, and learn from these unnatural disasters.

The landscape of evidence-based urban monitoring is primitive: right now, there is no well-established, affordable, high-resolution technique for urban monitoring. Existing technologies are far from ideal: physical sensors or cameras are expensive, susceptible to inclement weather, and take time to set up. Smartphone apps are commercially available yet suffer from poor reputations related to respecting privacy; nor are smart phones accessible for all socioeconomic groups.

Moreover, this is a universal problem: any city is at risk for responding poorly to various crises without a proper understanding of the behaviors of their residents. Urban monitoring technology could provide empirical evidence on human activity with which to assess the effectiveness of a city's social services, as well as help predict how a potential perturbation in the environment could achieve a more efficient, safe, and smart space. Additionally, every resident deserves the opportunity to receive up-to-date information about the human activity occurring around them so they can make an educated decision on how to proceed.

Our proposed solution involves cities investing in a cost-efficient and high-resolution method of urban monitoring through the use of Distributed Acoustic Sensing (DAS) fiber optic technology. Available telecommunication fiber networks, which are widespread in "smart cities" like New York and Singapore, can be cheaply converted into dense seismic sensor arrays that can detect surrounding ground vibrations. As a local example, we have access to data collected by the Penn State DAS Array, a 4km-long array snaking through campus and made up of thousands of sensors continuously recording. With our seismology and machine learning expertise, we intend to efficiently process and interpret the data into quantifiable outputs of human and vehicular activity.

Our solution offers a product that visualizes real-time human activity in its many forms: foot, bicycle, and car traffic. Using this tool, residents and cities alike can more intelligently navigate the myriad of challenges presented by natural disasters.

In the next section, we detail the potential use cases and benefits of this technology.

2 Use Case

There are three major use cases of our product that all help to alleviate the complexities of navigating, responding to, and learning from a disaster.

(1) **Empower individuals** to make informed decisions based on the human activity occurring around them;

(2) Help cities and relief efforts **track** those stranded throughout a disaster; and

(3) Allow cities to **learn from data** of human activity patterns during one disaster to better prepare for the next.

2.1 Empowering Individuals

Upon the announcement or discovery of a disaster, a resident of a smart city can go to our website and check out the real-time activity occurring around them. Our site presents a heatmap of human activity that can be distinguished into walking, biking, or driving. Our tool can answer questions like “Which streets are currently busy with traffic?”, “Where are people walking or running away from?”, or “Where are people congregating for safety?” Even without the ability to communicate with neighbors, police, or city officials, individuals will be able to gain information on the human activity in their area. It is for this reason we imagine our tool as the future “Google Maps of Human Activity”.

Had this tool been available during Hurricane Katrina, we imagine the greatest benefits for individuals to occur in the beginning: deciding how to exit the city, where local shelters were stationed, and what parts of the city were safe to navigate would all have helped residents get to safety as levees began to break; and each of these questions would have been possible to answer using our ideal tool.

2.2 Tracking

So long as there are people walking/driving on the ground, our detector will track their movements, and our AI `TraceNet` will be able to distinguish between the various activity types, as well as weather (using their unique frequency signatures). The resilience to surface weather and relatively low maintenance necessities make our service a great candidate for locating stranded individuals during a disaster.

Had this tool been available during Hurricane Katrina, the city of New Orleans would have had a useful tool with which to locate survivors across the city. This could only be possible in areas where people could walk on the ground, but many people were still stranded in their areas within New Orleans even after the waters from Hurricane Katrina significantly subsided [2].

2.3 Learning from the Past

With more time, more user testing, and a larger development team, we could offer a data analysis software feeding off of the same pool of data as `TraceNet` so that urban planners and city administrators can see how their city responds to disasters both large and small. This would empower cities to rework their disaster protocols. Cities would be able to investigate data from moments leading up to, during, and after a disaster to obtain a full picture of the human response to such a disaster in that city.

New Orleans would have benefited from this hindsight analysis tool after Hurricane Katrina. We suspect that our tool could help reveal improvements in exit routes, locations of shelters, and potential locations for significant flooding at which to rescue survivors first [2].

3 Development Timeline

Below is a rough timeline of major development milestones which we wish to achieve.

- **April - June**
  - **AI**
    - Add another category for the AI to predict: cars/vehicles
    - Develop a Graphic User Interface (GUI) program to more easily label DAS data for various activity types
    - Rework AI to handle multiple categories of activity (e.g., walking, bicycle, car, bus, etc.)
    - Perform a lot of manual labeling, possibly hiring another person to handle interpretation and labeling
    - Implement `TraceNet` using TensorFlow Serving
  - **Data Pipeline**
    - Settle on an efficient way to stream data from the DAS server to the web server
    - Develop automated methods to preprocess DAS data for AI learning and predictions
    - Secure GPU resources (Quadro RTX 8000) for deep learning training
  - **Website**
    - Investigate other softwares that could help handle the live functionality (especially Microsoft Azure, PowerBI, or ArcGIS)
    - Update website design
  - **Marketing & Branding**
    - Update logo
    - Contact urban planners within the PSU Urban Planning Department and in the industry to evaluate our product's use case further
    - Promote our brand on social media
    - Investigate other use cases related to urban planning, security monitoring, and investing
- **July - August**
  - **AI**
    - Finalize implementation of AI on live data
    - Explore the generalizations for multiple DAS arrays
    - Classify each activity further by level: low, medium, high, for example

- **Data Pipeline**
  * Debug and optimize the transfer of data to reduce latency
- **Website**
  * Host the website at a proper web server
  * Debug any and all connectivity problems should they arise
- **Marketing & Branding**
  * Provide marketable design inputs on the website
  * Beta test the product and collect user experiences to evaluate the product

◇ **August**

- Use this time as a buffer before the conclusion of the MVP Phase and submission on 10 August 2021.

◇ **Beyond the Nittany AI Competition 2021**

- Approach telecommunication companies to pitch our product to them and seek funding
- Test our product in a "smart city" (e.g., Philadelphia)
- Convert our work into a research paper
- Explore the various use cases beyond disaster relief identified throughout our MVP Phase

4 Technology

Our technical goal is to train an AI agent (which we have named `TraceNet`) to automatically identify human footstep signatures from DAS data in near-real time. Our general workflow can be summarized by the following:

(1) Preprocess the DAS data and manually label some human footstep signatures,

(2) Design an AI agent (`TraceNet`) and train it on the labeled data, and

(3) Use the trained AI to predict/isolate human footsteps on new DAS data.

4.1 DAS Data and Labeling

A fundamental principle of AI prediction is: the more training data that is available from which to learn, the more accurate future predictions can be. DAS data is basically a continuous stream of seismic data, collecting signals at a high frequency (250 Hz). The DAS sensors have been recording since April, 2019 until the present. Each sensor records 250 readings per second, and there are 2120 such sensors.

This wealth of information on which to train presents us with a different challenge: labeling all of that data. Certainly we cannot label all of it, nor even a majority. However, for the purposes of training our prototype model, we selected a portion of sensors (175 sensors out of the 2120) from a walking-only path at the heart of campus (see Figure 1). We choose this path because it best represents student activities on campus without other confounding

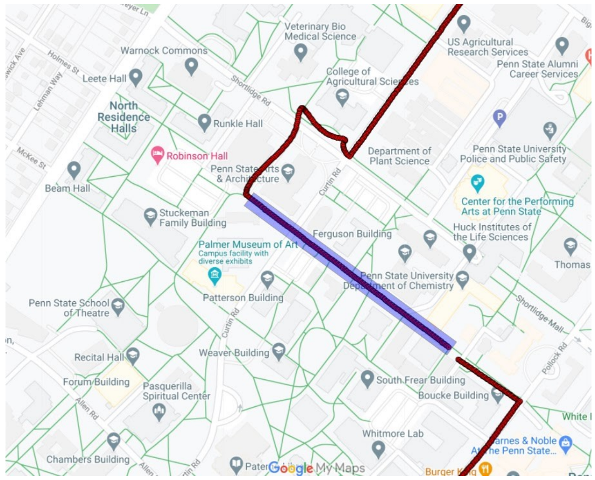

Figure 1: The segment (highlighted in blue) of the DAS array most convenient for labeling walking data. This portion of the array includes 175 sensors.

activities. Still, even after selecting this limited segment, the amount of data recorded continuously since April 2019 is too much to handle. Instead, we developed the following data sampling method that could capture a representative snapshot of daily walking activities:

(1) Partition each day into five time segments to capture different levels of walking activity:

- ◇ **Morning**: 7AM - 12PM
- ◇ **Afternoon**: 12PM - 5PM
- ◇ **Evening**: 5PM - 9PM
- ◇ **Night**: 9PM - 12AM
- ◇ **Early Hours**: 12AM - 7AM

(2) Uniformly randomly select 65 segments in each category, totalling to 360 segments, which is equivalent to 15 days of data.

This random sampling process produces a robust representation of all kinds of walking patterns ever since April, 2019.

We performed data preprocessing to prepare the data for labeling. The preprocessing steps are:

(1) Separate the raw data into hourly files. Each hourly file can be treated as an image.

(2) Downsample each hourly image from 250 Hz to 10 Hz.

(3) Normalize on both axes.

(4) Separate each hourly image into twelve sets of five minute images.

- ◇ Each five minute image will have dimensions 175 px × 3000 px (175 is the number of sensors, 3000 is from 300 seconds × 10 Hz = 3000 readings within 5 minutes)

(5) Manually label human footsteps.

In total, we preprocessed and labeled 4320 images, which took about 50 hours of work. Figure 2 shows an example of data and its labeled image. The labeled image is essentially binary, where 0 corresponds to no walking activity, and 1 means there is walking activity.

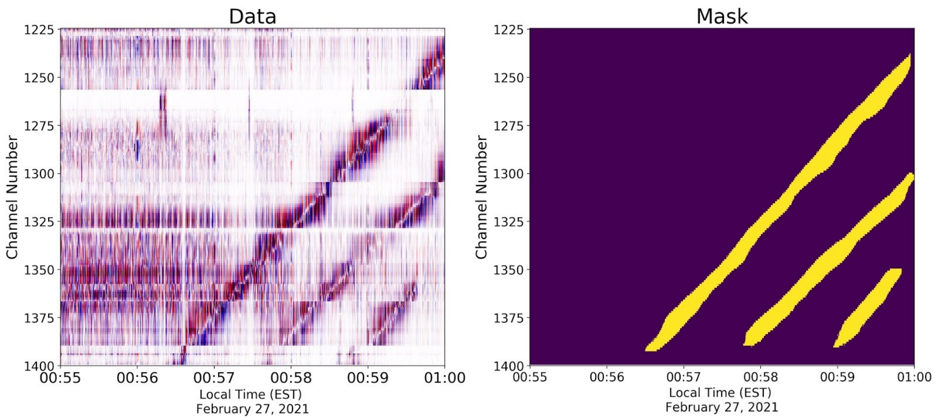

Figure 2: Left: An example of Channel Numbers vs. Time raw data over a 5 minute span. Right: our manual labeling mask to identify human footsteps from the raw data on the left.

4.2 AI Design — `TraceNet`

Via the preprocessing procedure described above, the problem of classifying human foot traffic boils down to an image segmentation problem. For that reason, we developed a Convolutional Neural Network (CNN), `TraceNet`, to perform this classification. CNN is well known for its ability to locate objects as instructed during training. For our prototype, we designed `TraceNet` roughly based on U-Net [3], which was a CNN designed to classify tumors in medical images. The architecture underlying `TraceNet` is illustrated in Figure 3. Figure 4 details the corresponding blocks that were used in Figure 3.

At its deepest layer (i.e., `Conv6` and `Deconv1`), `TraceNet` transforms input data into 1×1 representational features; these features allow the AI to learn intrinsic patterns between the input and output images, boosting its accuracy.

We split our dataset into 95% (4104 images) for training, and 5% (216 images) for validation. The deep learning training is set to stop when the validation loss convergences to a very small threshold. The training took approximately 3 hours to complete.

4.3 Prediction on New Data

We applied our AI, `TraceNet`, on selected data from 1 March 2021 to 13 March 2021. Next, we downsampled the predicted images by a factor of 0.1 along the time axis, making the new dimensions for each image 175 px $\times$ 300 px. These dimensions correspond to the 175 sensors (each having distinct GPS coordinates), and 300 seconds of data. At each point of location and time, the predicted labels have a value of 0 or 1: 0 meaning "no walking activity", and 1 meaning "there is someone walking".

Finally, we converted the predictions into a readable CSV format with the intention of visualizing the predictions on our website. Figure 5 shows a snippet of such a CSV file.

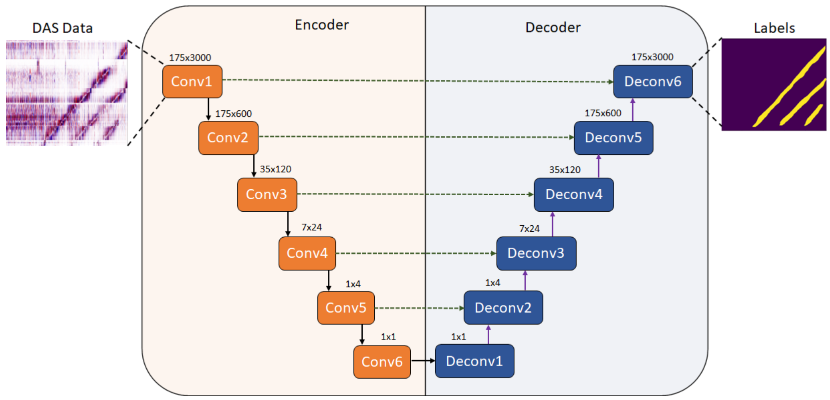

Figure 3: A visualization of our deep learning architecture for `TraceNet`. The model takes in DAS data images as input, and outputs the predicted labels. Each `Conv` block represents the composition of operations illustrated in Figure 4(a); whereas each `Deconv` block corresponds to Figure 4(b).

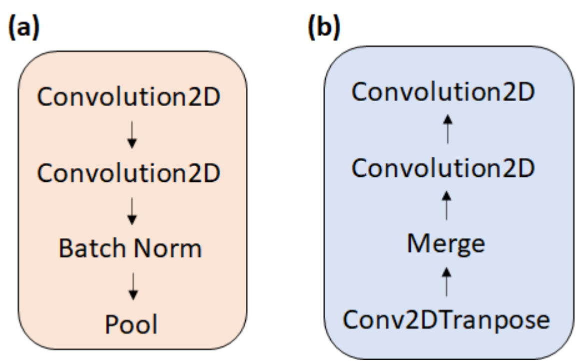

Figure 4: (a): A `Conv` block within the deep learning architecture (see Figure 3). Each `Conv` block consists of two convolutional layers, a batch normalization layer, and a pooling layer. (b): A `Deconv` block. Each `Deconv` block consists of a `Conv2DTranspose` layer which acts as an unsampling layer; a merge layer which concatenates blocks from `Conv` and `Deconv` at respective layers; and two convolutional layers. The convolutional layers use ReLU as activation functions, and padding is set to "same".

5 Data Sources

Running beneath Penn State University Park's campus is a fiber optic array which uses Distributed Acoustic Sensing (DAS) technology to continuously measure vibrations occurring in the vicinity of the line; the DAS array is part of the Fiber Optic foR Environmental SEnsEing (FORESEE) project at Penn State. Our continuous data comes from this detector.

More specifically, we make use of otherwise unused underground telecommunication fiber optic cables buried in a conduit at a depth of 1 meter. The DAS array makes continuous strain rate measurements at 500 Hz sampling frequency with a 10 meter gauge length and 2 meter channel spacing. In total, there are 2120 sensors across Penn State's campus along this line. All data is currently stored within the College of EMS's server. The data will be

	channel_number	Date-Time	Latitude	Longitude	Activity
5500	1300	2021-03-05 10:00:31	40.800747	-77.863889	0.0
5501	1301	2021-03-05 10:00:31	40.800736	-77.863870	1.0
5502	1302	2021-03-05 10:00:31	40.800725	-77.863850	1.0
5503	1303	2021-03-05 10:00:31	40.800714	-77.863831	1.0
5504	1304	2021-03-05 10:00:31	40.800704	-77.863812	1.0
...	...	...	...	...	...
5695	1320	2021-03-05 10:00:32	40.800530	-77.863505	0.0
5696	1321	2021-03-05 10:00:32	40.800519	-77.863486	0.0
5697	1322	2021-03-05 10:00:32	40.800508	-77.863466	0.0
5698	1323	2021-03-05 10:00:32	40.800497	-77.863447	0.0
5699	1324	2021-03-05 10:00:32	40.800486	-77.863428	0.0

Figure 5: An example CSV file that is used by the website to visualize human activity. This CSV file was built from the output of our activity categorizing AI `TraceNet`.

released via Penn State Commons after 2021.

As described in §4.1 and illustrated in Figure 1, our prototype focuses on data recorded by Channels Ch 1225-1400 along a pedestrian-only path. This fiber also goes across Curtin Road (at Ch 1267). To reiterate the preprocessing: data were first downsampled to 10 Hz to reduce the computation and storage cost; then a high-pass filter ($>$ 1 Hz) was applied to remove low-frequency noise. We chose 15 random days from January, 2021 to early March, 2021, including part of Penn State's winter break, as well as weekdays and weekends.

DAS technology is not intended to identify individuals, meaning our data is anonymous and ambivalent to "who?" Instead, this detector asks "whether" (i.e., whether there is presence or absence of activity), while our AI `TraceNet` attempts to answer "how?" (i.e., how that activity was generated).

6 Team Capabilities

To meet the challenges of implementing our solution, our team consists of five members with specialized skill sets appropriate for their role in the project. In order of their role in the workflow pipeline, our team members are Zi Xian Leong, Junzhu Shen, Ian Lee, Robert Krimetz, and Raymond Friend.

- **Zi Xian Leong (Team Lead)** is a PhD student in the Department of Geosciences working on deep learning applications in Geophysics, and he has been developing the AI framework for `TraceNet`. Zi Xian produced our team timeline to keep our group on track to hit certain milestones before the conclusion of the Prototype Phase, as well as beyond.
- **Junzhu Shen (Technical Lead)** is a PhD candidate in the Department of Geosciences currently working with the Penn State DAS Array data as part of his research, and he provides our team access to DAS data as well as knowledge on best practices for processing and filtering the raw DAS data. He also provides guidance on results interpretation.

- **Ian Lee (Technical Lead)** is also a PhD student in the Department of Geosciences with a focus on Cryoseismology and Data Science. He has been developing the interactive map visualizations for the output and setting up an efficient data transfer protocol between the back end and front end. Ian worked along with Junzhu to label all 4320 images.

- **Robert Krimetz (Design Lead)** is an undergraduate Data Science student with experience in data analytics. He has been developing our full stack Django Web Application to store and visualize the analyzed DAS data. Robert works closely with Ian on the front end of our Trace website.

- **Raymond Friend (Marketing Lead)** is a PhD student in the Department of Mathematics with a focus in Logic. As our Marketing Lead, he made use of his experience in video production through the Adobe suite (e.g., Premiere Pro) to create our prototype video; and his experience in LaTeXto produce this Prototype Phase Document. He works with Zi Xian to set the strategic direction for our product.

7 User Interface

At the heart of our user experience is the ability for individuals to track changes in human activity (including footsteps, vehicles, bikes) in their area. In order to achieve this goal, we developed a prototype Web Application to visualize the outputs of `TraceNet`. Currently, users can interact with a Google Maps interface that uses heatmap plotting to illustrate the level of human activity along the DAS array over that portion of time. See Figure 6 for a screenshot of what a typical user would see using our prototype Web Application.

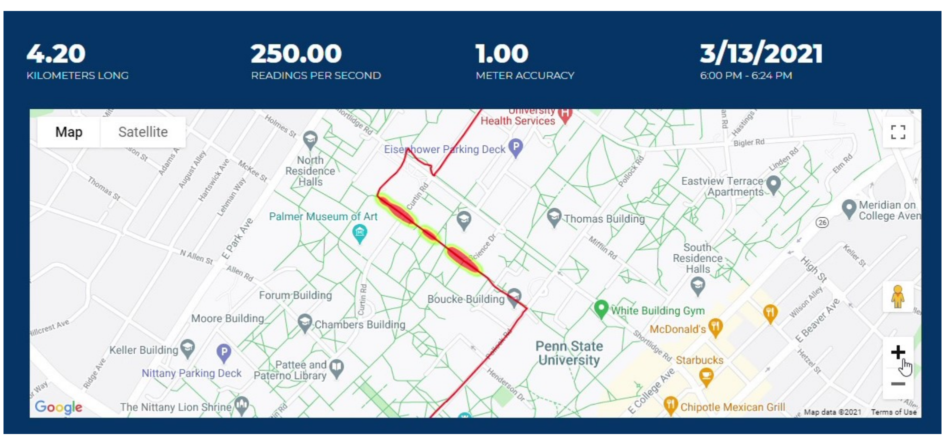

Figure 6: A screenshot of our website, showcasing the Google Maps interface that displays the output of `TraceNet` over a particular time span as a heatmap. This is along Channels Ch 1225-1400 (a pedestrian-only path) on 13 March 2021 during 6:00PM-6:24PM.

The website is a full stack Django application which implements an SQLite Database as a back end. The front end is written in CSS, HTML, and JavaScript. The interactive Google Map interface is written using the Google Cloud API and Dynamic JavaScript. A batch manager written in Python is used to connect the cleaned CSV output from `TraceNet` to the Web Application's database.

References

[1] Hannah Ritchie and Max Roser. Our World in Data. 'Natural Disasters' (2019, November). Retrieved March 28, 2021, from https://ourworldindata.org/natural-disasters.

[2] PBS. '14 Days — a Timeline' (2005, November 22). Retrieved March 28, 2021, from https://www.pbs.org/wgbh/pages/frontline/storm/etc/cron.html.

[3] Ronneberger, Olaf, Philipp Fischer, and Thomas Brox. 2015. 'U-Net: Convolutional Networks for Biomedical Image Segmentation'. ArXiv:1505.04597 [Cs], May. http://arxiv.org/abs/1505.04597.

Trace