Predicting a Meteotsunami with Recurrent Neural Network

Anonymous Author(s)
Affiliation
Address
email

Abstract

In this project I am attempting to predict meteotsunamis - long ocean waves caused by rapidly changing atmospheric pressure. Meteotsunamis have very similar properties to famous and dangerous tsunamis, but there is much less research going on them: no one tried to use a large dataset (covering a few years of sea observations at a specific location) to train some supervised-learning algorithm on it. I applied Recurrent Neural Network (RNN) to such a dataset and somewhat succeeded in predicting a meteotsunami given the atmospheric pressure at the previous timesteps. However, the model built by RNN turned out not to be very accurate.

1 Introduction

In this paper I am dealing with so-called meteotsunamis. In short, meteotsunami is an ocean wave which has very similar properties to a tsunami wave, but is caused by atmoshperic pressure change as opposed to the movement of the ocean floor. Below I am going to describe all these terms in more detail.

1.1 Tsunami

A tsunami is an ocean wave, which differs from an ordinary wind wave in that it has a very long length and period. Essentially, a tsunami can be thought of as a very fast tide.

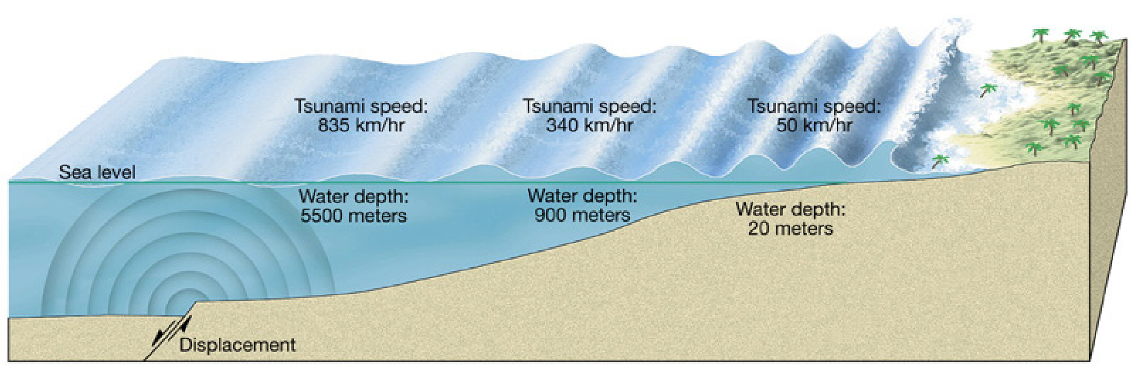

Figure 1: Tsunami formed by the ocean floor displacement. (From [4]).

Tsunamis are usually formed by movement of the ocean floor (earthquakes, volcanos etc. - see Figure 1) - when it moves even a little bit, the whole mass of ocean water starts to adjust to it, causing a wave. Such a wave may not look very high and dangerous initially, but it carries an enormous amount of water mass which is much larger than the mass carried by a regular wind wave of the same height (because a tsunami is a long wave).

That said, some tsunamis have a very small height (a few centimeters), and thus are not noticeable by people living on the coast. However, if a tsunami has a height of more than a meter, it can cause great damage to the shore. No surprise then that all those colossal walls of water of height more than 10m, which people usually associate the term ”tsunami” with, literally devastate the coastline (Figure 2).

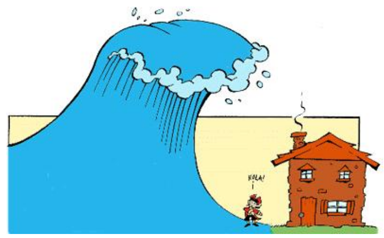

Figure 2: Possible consequences of an earthquake. (From [6]).

1.2 Meteotsunami

A meteotsunami is a wave with the same properties as a tsunami (i.e. long period and large wavelength), but which is formed in a different way. As follows from its name, a meteotsunami has an atmospheric origin - in this case, the role of moving ocean floor is played by rapidly changing atmospheric pressure - e.g. which happens when a storm front is passing by (Figure 3).

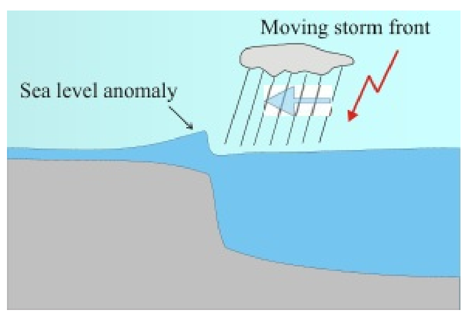

Figure 3: A meteotsunami is being formed by a passing storm front. (From [3]).

You might wonder why meteotsunamis are not that well known given how many storms are going around. The catch here is that the change in atmospheric pressure has much less of an impact on the ocean than an earthquake or a volcano. Thus, generally, meteotsunamis are not as dangerous as tsunamis (which draws less attention to them), and several conditions need to be satisfied for them to be formed [2]:

- There should occur a meteorological event causing a rapid change of atmospreheric pressure above the sea
- The speed of this event should match the speed of waves currently present at deep water (causing resonance)
- The harbor should have a shape that amplifies the long wave coming to it

The first two conditions imply that meteotsunamis are quite rare, while the third one adds that meteotsunamis should occur mostly at some specific georgaphic locations. And, indeed, a meteotsunami is a relatively well-known phenomenon in Spain’s Balearic Islands (there it is called *rissaga*), Japan (with a local name *abiki*), and other Mediterranean countries like Malta or Croatia.

It should be mentioned that although most of meteotsunamis are small and not dangerous, sometimes they can be as tall as 4 meters [2], in which case they are hard to overlook (Figure 4).

Figure 4: ”Rissaga” at Menorca Island (Spain) on June 15th, 2006. The damage was estimated to be about 30 million euros. (From [2]).

2 Problem Description

While there is a tremendous amound of research going on about predicting and analyzing tsunami waves, meteotsumamis get much less attention (apparently, it happens because meteotsunamis are localized and less dangerous, and thus researchers mostly focus on much more destructive and global tsunamis). You will also not be very surprised to learn that main research on meteotsunamis is going at the places where it poses a real threat - like Balearic Islands and Croatia.

The main problem related to any kind of long waves that researchers solve is, of course, alerting the coast that a (meteo)tsunami is coming. The earlier the alert comes, the better chances people have to save their lives and properties before the wave arrives. This makes the *prediction* of these waves a crucial task.

As we all know, the level of science novadays does not allow us to predict catastrophic events like earthquakes (and, hence we cannot predict a tsunami following it either). However, meteotsunamis have an atmospheric origin, and thus there is a possibility that we can predict them given the data about atmospheric pressure nearby. It should be noted here that, although there exist lots of models explaining how meteotsunamis appear, no research tried to use large empirical data (covering a few years of observations) to predict meteotsunamis so far.

So, in this project I am trying to apply a supervised learning algorithm to the empirical data about the atmospheric pressure and sea level in an attempt to predict the occurence of a meteotsunami.

3 Data

As was mentioned above, in this project I am trying to predict meteotsunami given the atmospheric pressure data. So, what could we use as an *output* for our algorithm? How can we ”measure” a meteotsunami? A standard way to measure long waves is by measuring *variance of the sea level averaged over time*. Indeed, variance in this case describes the deviation of the sea level from normal (over time), which intuitively seems to be a good indicator of how ”wavy” our ocean is. Note, that we average over a large time period (an hour) in order to capture long waves.

So, in short, the output data is the variance of the sea level averaged over one hour. I used measurements of a specific detector at Vancouver Island’s coast for 3 years (2007-2009), which yields over 20000 output data samples after averaging (each sample is just a *single real number*). This data is available online at [7].

It should be noted that around Vancouver Island only 3 meteotsunamis were detected in 2007-2009 - one in each of the years [5], and our output data nicely reflect that - there is a clear spike in the graph corresponding to a meteotsunami happened (Figure 5), which will be later referred to as a ”meteotsunami spike”.

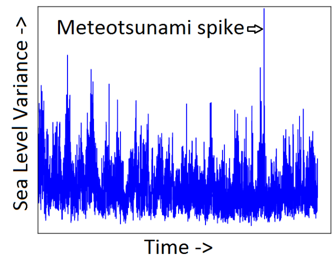

Figure 5: The plot of the output data for the year 2009 versus time - the largest spike corresponds to a meteotsunami happened in the fall of 2009. Noteworthy, almost nobody saw this and other meteotsunamis at Vancouver Island - they are noticeable only at some very specific areas. Those who did see it, though, were quite impressed - for instance in 2008 a fisherman (who was the only human being at that location at that moment) said that he never imagined the ocean behaving like this [5].

What about our *input* data? In this case we can use measurements of a barometer located nearby (more precisely, near Deep Cove Elementary School) for the same time period (2007-2009). In the same manner as with our output data, we average these barometer measurements over one-hour period, thus getting the indication of how fluctuating atmospheric pressure was at that moment. The data is available online at [8].

It is worth mentioning that even a few years ago such data from British Columbia coast would not have been available (barometers did not give us minute-by-minute measurements), making this project impossible to accomplish (at least on data from BC). But professor Andrew Weaver initiated the project which can be described as ”give a barometer to every elementary school” [8], and made the measurements done by those new barometers publicly available. Thus, in a sense, elementary schools move the ocean science forward in British Columbia.

Getting back to my project, here I can reiterate that the data I used for it is two real-valued temporal series (from 2007 to 2009): input (atmospheric pressure variance) and output (sea level variance at the *next timestep*). The problem of predicting a meteotsunami in this case reduces to fitting this data by applying some supervised learning algorithm to it.

It should also be clarified here that using the next timestep for the output relative to the input is necessary for ensuring that no model has access to the atmospheric pressure measurements while trying to predict the sea level for the same timestep (otherwise, it would not really be a prediction).

4 Algorithm - Recurrent Neural Network

Although at the first look the problem may seem to be extremely simple - one input, one output - the catch here is that the sea level variance (output) depends not only on the corresponding atmospheric pressure at the previous timestep (input), but also on the atmospheric pressure and sea level from the *earlier* timesteps.

Therefore, if we want to predict a meteotsunami as accurate as the data allows, we might want to use the algorithm which can potentially capture these temporal dependencies. Note, that the most popular supervised learning algorithms like linear regression do not seem to be promising as they do not take the temporal component (i.e. order) into account.

One of the algorithms that *take* the temporal component into account is so-called Recurrent Neural Network (RNN). RNN differs from a classic feed-forward Neural Network (FFNN) in that it has connections from hidden units to themselves (Figure 6). This architecture allows the output to depend on the dynamics of the network at previous timesteps through these hidden-to-hidden connections.

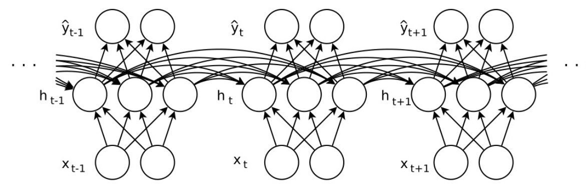

Figure 6: A typical architecture of Recurrent Neural Network (RNN). Here RNN is depicted as series of feed-forward neural networks, with hidden units at a given timestep connected to hidden units at the next timestep. (From [1]).

Previously RNN were not widely used for solving problems with temporal dependencies, because they are very hard to train with standard algorithms like backpropagation. However, recently Martens and Sutskever (2010) applied a method called "Hessian-Free optimization"(HF) to train RNN, and showed that in combination with powerful GPU units it allows to train RNN much faster and better (allowing it to capture much more complex temporal dependencies than before) [1].

I modified their code available online to make it work for my particular dataset. I used RNN with 50 hidden units with *tanh* activation function, 1 input unit and 1 output unit (with identity activation function). I used default parameters (provided by Ilya Sutskever's code) for the initialization of weights and for HF training algorithm, since he found them to work moderately well on all datasets he tested [1].

I used the first 2/3 of the data available (roughly years 2007-2008) as training data - it was standardized and divided onto 100 possibly overlapping sequences of 100 timesteps each starting from a random point in time. The RNN was then trained on these so-called minibatches. The rest of the original data (roughly the year 2009) was used for testing the results.

Note, that I also used basic linear regression algorithm for comparing its performance to RNN and validating that the recurrent neural network can do a better job.

5 Summary of Results

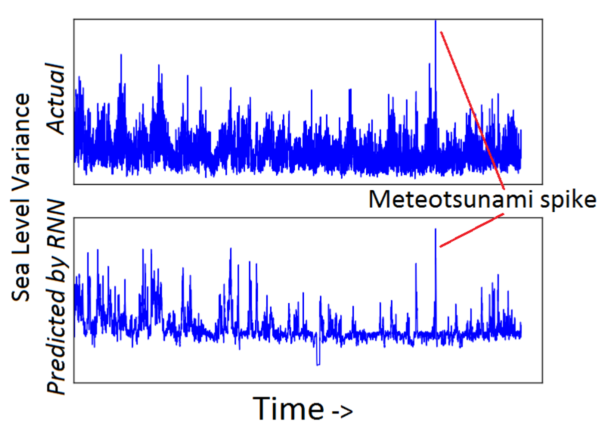

Figure 7: Top: the plot of the output data for the year 2009 versus time - the largest spike corresponds to a meteotsunami happened in the fall of 2009. Bottom: the plot of the *predicted* output for the year 2009 versus time - note, that although the prediction is quite bad overall, RNN is somewhat able to track the meteotsunami spike.

After training the model, I found that it was able to predict the meteotsunami spike exactly at the right time on the test dataset (Figure 7). However, note that the predicted spike on the graph is not as prominent compared to the actual output - RNN was not able to fit the rest of the output.

Also, not surprisingly, the linear regression model was unable to predict the meteotsunami happening.

Unfortunately because of the time constraints of the project (and also since I did not have any Linux machine with GPU available for running the code), I was unable to try all the different parameters (e.g. number of hidden units used) and see how the network performs.

6 Conclusions and Discussion

In this project I was able to train RNN to predict sea level variance spike corresponding to a meteotsunami happened at Vancouver Island in the fall of 2009, given atmospheric pressure and its dynamics at previous timesteps.

Although this might seem to be a positive result, the problem is that RNN was completely unable to fit the data overall, showing that its representation of the problem is not very accurate - and thus it is likely to fail on a different dataset (in this project I was literally predicting a *single* meteotsunami). In addition, I did not compare the performance of RNN to the performance of other advanced algorithms like Gaussian Processes for temporal series - maybe they would do a much better job at capturing temporal dependencies in the data.

It can be the case that playing with parameters of RNN could be beneficial, as well as adding to the network connections from the output units back to the hidden units. It may be that these modifications will improve the results.

However, it may also be that problems with fitting the data are related not to the algorithm applied, but to the data itself. I used atmospheric pressure measurements only at one location - it is quite possible that using the measurements from several nearby locations could improve the performance of the RNN dramatically.

Finally, I want to point out that even if RNN will not be proven to be efficient for solving the problem of a meteotsunami prediction, this network seems to be suitable for at least *initial* analysis of such problems with temporal dependencies. RNN is relatively simple, easily customizable for any type of input-output data, quite powerful and fast (with the availability of GPUs and a good training algorithm like HF optimization).

I believe that these properties make RNN a perfect candidate for a generic black-box algorithm that a researcher might want to apply to a new dataset to understand its dependencies better, before choosing the best model available.

References

[1] Martens, J. & Sutskever, I. (2011) Learning Recurrent Neural Networks with Hessian-Free Optimization. *Proceedings of the 28th International Conference on Machine Learning*

[2] Monserrat, S., Vilibic, I. & Rabinovich, A.B. (2006) Meteotsunamis: Atmoshperically Induced Destructive Ocean Waves in the Tsunami Frequency Band. *Natural Hazards and Earth System Sciences* **6**: 1035-1051.

[3] Rabinovich A.B. et. al., Meteorological Tsunamis: Destructive Atmosphere-Induced Waves Observed in the World Ocean. Web. 30 Nov 2011.

[4] Schott R. (2011), GSCI 340 - Environmental Geology Lecture 11. *Fort Hays State University*. Web. 1 Dec 2011.

[5] Thomson, R.E. et. al. (2009) Meteorological Tsunamis on the Coasts of British Columbia and Washington. *Physics and Chemistry of the Earth* **34**: 938-955.

[6] Official website of Department of Cooperative Governance, Republic of South Africa. Web. 1 Dec 2011.

[7] Official website of Department of Fisheries and Oceans, Canada. Web. 25 Nov 2011.

[8] Official website of School-Based Weather Station Network, Greater Victoria. Web. 27 Nov 2011.