An Analysis of the July 22, 2021 Heatwave at Shannon Airport

Davis Sandefur
12250836

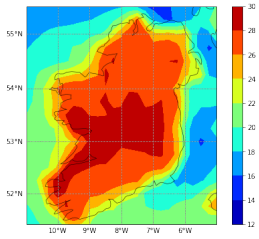

19 April, 2022

1 Introduction

In the last two weeks of July 2021, the island of Ireland experienced another heat wave, with temperatures reaching far outside their normal summer means. Temperatures during the heatwave were consistently over 20°C, with many days having highs over 25°C. These numbers are well outside Ireland's average temperatures and highs for the month of July. At Shannon Airport, which will feature in this report, the highest temperature was reported at 29.3°C at 15:00 UTC on 22 July, 2021. To examine this heatwave, we will first run multiple forecasts using the WRF numerical weather prediction model. These forecasts will be run for the hottest day of the year at Shannon Airport, July 22, 2021, and the forecasts will be run starting at seven, three and one day in advance of the event, as well as a forecast run on the same day. This data will be compared to the reanalyzed data available via the ERA5 datastore. We will then more closely examine the hottest day of the year at Shannon Airport, which had the highest recorded temperature of long-running Met Éireann weather stations. We will do so by comparing the actual values reported by Met Éireann to climatological data obtained from their data. This will allow us to see how far outside the norm this heat wave was. Finally, multiple skill scores will be run on our forecast models, comparing them both to climatological prediction data and persistence prediction data, giving an overview of how accurate our models were and whether they were good indicators of the event.

2 Data and Methods

The data comes from three major sources. The first source is the Met Éireann historical data archive. This archive supplies historical observation data from the various stations across Ireland, available in either an hourly, daily or monthly basis. We will use the hourly data from station number 518, based at Shannon Airport. The hourly data contains hourly observational records of temperature going back to January 1, 1992 and continuing on through early 2022. Thus the observation data for our even is included. This will slightly bias our results when we compare the 2021 heatwave to climatological data, as we should ideally have 30 years of data ending at the decade before the event (in this case, 1991-2020); thus, when we use the 30 years worth of data to get climatological information, our results will include our heatwave

event, meaning that the quantiles will possibly be higher than they would have been otherwise.

The next source of data is the ERA5 hourly reanalysed data from the EU's Copernicus. This data combines observations with physical laws to give a reanalysed dataset of the weather stretching back to 1979. This dataset will allow us to get a reanalysed look at the temperatures across the entirety of Ireland during the heatwave of July 2021, particularly at 17:00 UTC on the day of the peak. This will then be used to compare to the forecast models to see how the models differed from reality.

The last dataset used was the NCEP GFS Global Forecast Grids. These grids provide data in three hour intervals that is used to run forecasts. The data contains the conditions around the boundary set by the user, in our case a large area encompassing Ireland. These are then used by WRF to generate a forecast model over a given time range.

Once the data was all gathered and processed, the WRF model was run. The WRF model that was used was centered at the latitude of 54.8° and a longitude of −7°. There resolution was set to 10km and there were 120 grids in the east-west direction and 110 grids in the south-north direction. This was done to maximize the amount of boundaries over water, though some of north-eastern France did have to be included. The governing equations were solved every minute, and the data was output after every hour, matching it to the hourly data received from Met Éireann and the ERA5 data. All forecasts ended at 00:00 on July 23, 2021, allowing us to capture the full day of the event; they started at 00:00 on the 22nd, 21st, 19th and 15th of July for the same-day, one-day, three-day and seven-day forecasts, respectively.

3 Results

Figure 1 below shows the temperature from the ERA5 data for Ireland at 16:00 UTC on July 22, 2021. This is the time the highest temperature was recorded by Met Éireann at the Shannon Airport station. As can be seen, temperatures in Ireland were high at the time, with only a little swath of the island having temperatures below 20°C and the vast majority of the island having temperatures above 26°C. We can see that Shannon Airport falls within the zone that recorded a temperature of over 28°C, which shows that the ERA5 data does correlate with the actual data provided by the weather station there.

Figure 1: ERA5 Temperatures (°C) for Ireland at 15:00 UTC on July 22, 2021

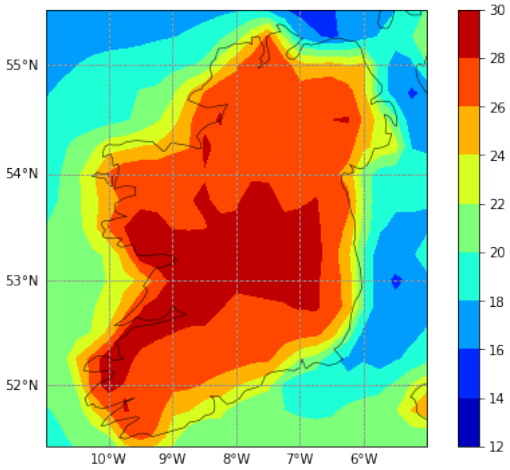

Next, we will look at how this temperature compares to climatological data in Ireland. To do this, thirty years' worth of data was collected from the Met Éireann station at Shannon Airport. The data for July across the thirty years was then analyzed and split into quantiles based on hour of the day. This was then plotted against along with the data given by the Met Éireann station for July 22. As can be seen in Figure 2, below, the temperatures on this day were well outside climatological norms, eclipsing even the 99th quantile of historical temperatures for most the day.

Figure 2: July 22, 2021 compared to climate data for July at Shannon Airport

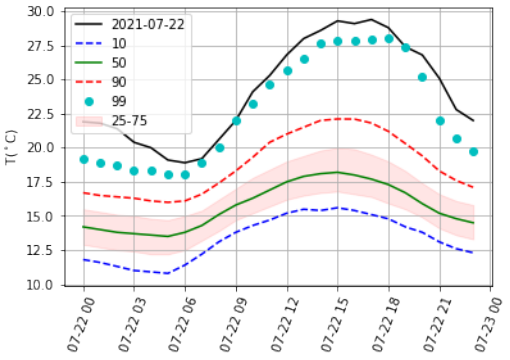

Now we turn to the forecasts. First, we will look at the four forecasts at the time of maximum temperature to get an idea of how they varied as the event approached. These results can be seen in Figures 3. We can see from this that the seven day forecast predicted that only one place in Ireland would get above 26°C, with most the country having temperatures between there and 20 degrees. By the time of the three-day forecast, the predictions had changed dramatically, with most the country predicted to be above 26°C. The one-day and same-day forecasts did not lower the temperatures in any area, but expanded the area that would experience higher temperatures, especially the area that would experience the most extreme temperatures

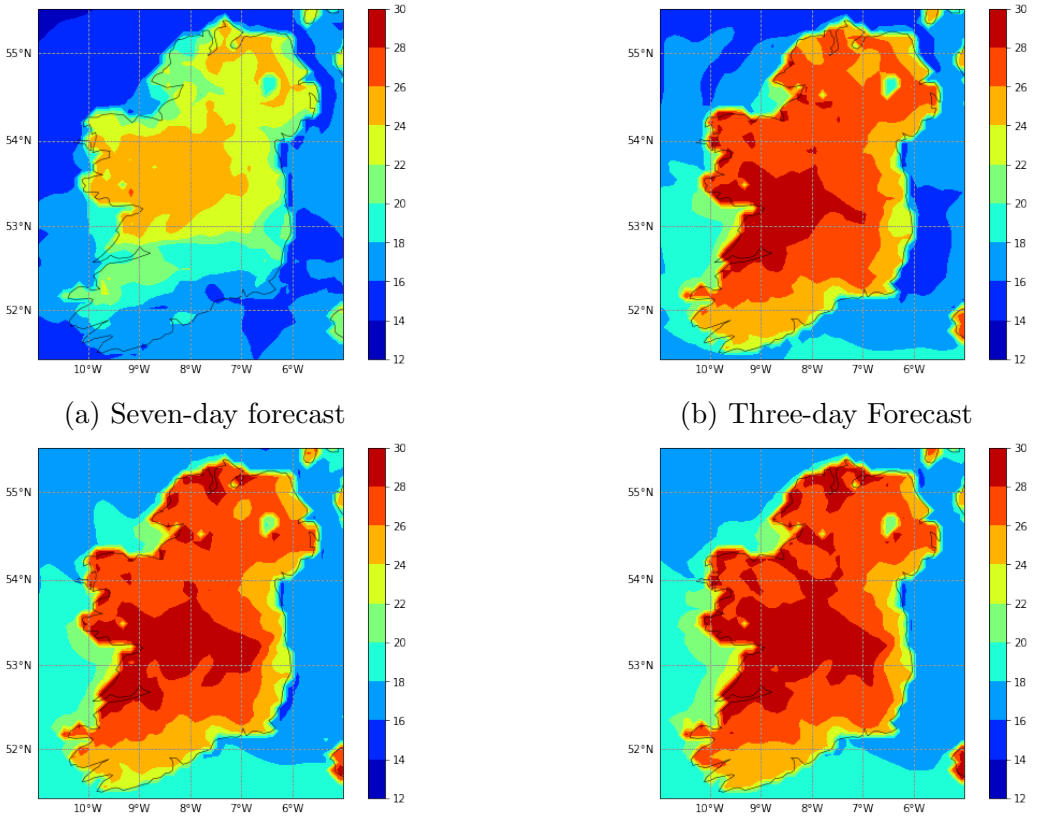

(a) Seven-day forecast

(b) Three-day Forecast

(c) One-day Forecast

(d) Same-day Forecast

Figure 3: Temperature on July 22, 2021 at 15:00 UTC as predicted by forecasts

A quick glance at the four figures in Figure 3 compared to the historical

data from ERA5 in Figure 1 shows that there were some differences in predictions, with the forecasts predicting much higher temperatures than did appear in some places. We will examine this directly for Shannon Airport now, looking at how temperatures evolved over the day in the forecasts as well as in real life and compared to average climatological models. Figure 4 shows how the forecasts fared with the reported temperatures from the Shannon Airport station, as well as the historical climate mean and the persistence (temperatures from July 21, 2021) data at the same station. These results can be seen in Figure 4.

Figure 4: Forecasts Compared to Observed Data, Historical Averages and Persistence Data

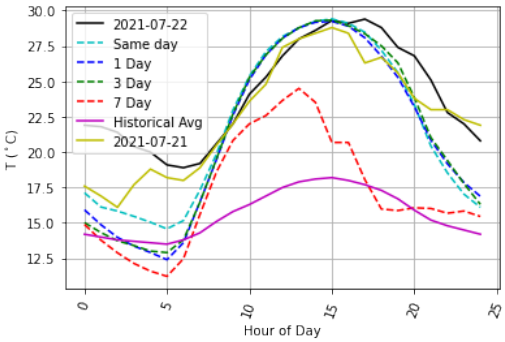

Overall, it looks like the forecasts predicted the maximum temperature with fairly good accuracy, though they did not do as well in predicting the temperature at other parts of the day, particularly overnight and afternoon, though they were reasonably good in the morning even if they over-predicted the temperature during these periods.

To quantify how accurate the various means of prediction are, we will use a variety of skill scores. The first skill score is simply the mean bias. It will be determined according to (1), where f_i is the forecast value at a specific hour and o_i is the observed value. Thus, if the forecast under-predicts the temperature, the bias will be negative. The results of this calculation can be seen in Table 1.

$$\sum_{i=1}^{25} \frac{f_i - o_i}{25} \tag{1}$$

Prediction Method	Bias
Persistence	-1.31
Climatology	-8.34
Same-Day Forecast	-2.25
One-Day Forecast	-2.79
Three-Day Forecast	-2.71
Seven-Day Forecast	-6.76

Table 1: Bias of Prediction Methods

We can see that, on average, all prediction methods undershoot the actual recorded temperatures for the day, with the persistence measure performing the best and climate performing the worse. It can also be noted that the seven-day forecast performs closer to the climate measure than to any of the other measures.

The next skill score to be used is the average Root Mean Square Error (RMSE) according to (2) across all hours of the day, where t_p and t_o are the predicted and observed temperatures, respectively. We average this across 25 total hours, including 00:00 on July 23 as our final point. The results can be seen in Table 2 below.

$$\sum_{i=1}^{25} \sqrt{\frac{(t_p - t_o)^2}{25}} \tag{2}$$

Prediction Method	RMSE
Persistence	2.14
Climatology	8.61
Same-Day Forecast	3.39
One-Day Forecast	4.06
Three-Day Forecast	4.09
Seven-Day Forecast	7.46

Table 2: RMSE of Prediction Methods

Once again, climatology has the highest error, with the seven-day forecast having the second highest. Persistence is again the best predictor used, with the same-day forecast being the second best. The one-day forecast and three-day forecasts are quite similar with roughly equal error between them.

We will also use the Anomaly Correlation Coefficient (ACC) as described in (3). This skill score directly compares the observations and the forecast by using the climate as our baseline. Because of this, the score when compared against climatology is meaningless and will not be considered. In the following equation f represents the forecast values, o the observed values and c the climatology values. The results can be seen in Table 3.

$$ACC = \frac{\overline{(f-c)(o-c)}}{\sqrt{\overline{(f-c)^2} \cdot \overline{(o-c)^2}}} \tag{3}$$

Prediction Method	ACC
Persistence	0.97
Same-Day Forecast	0.92
One-Day Forecast	0.88
Three-Day Forecast	0.88
Seven-Day Forecast	0.52

Table 3: ACC of Prediction Methods

Yet again, it shows that persistence is the best forecaster, even ahead of the same-day forecast. The one-day and three-day forecasts had an equal ACC score out to two decimals, while the seven-day forecast lagged quite a bit behind, barely providing any predictive power over random chance.

The final skill score we will look at is the V score. This score compares the forecast directly to climate as opposed to removing the climate values from the prediction as the ACC score does. Thus each forecast will give a number between zero and one, with zero being a forecast that predicts climatological average for every hour and one being a perfect forecast. The V score is calculated with (4) where f_c is climate forecast and f_o is observed and f_f is forecast prediction and the results are found in Table 4.

$$V = \frac{1}{25} \sum_{i=1}^{25} \frac{f_{ci} - f_{fi}}{f_{ci} - f_{oi}} \tag{4}$$

Prediction Method	V
Persistence	0.84
Same-Day Forecast	0.68
One-Day Forecast	0.60
Three-Day Forecast	0.61
Seven-Day Forecast	0.18

Table 4: V Score of Prediction Methods

As is expected from the other skill scores, persistence forecasting is the best predictor under the V score as well, with same-day coming in a distant second. The one-day and three-day forecasts were roughly equal, while the seven-day forecast was closer to zero (climatological predictions) than any of the other forecasts or observed temperatures.

4 Discussion

The first thing that can be noticed from our results was that the temperatures on July 22, 2021 were extremely high. A good swath of the country had temperatures over 28°C at 15:00 UTC (16:00 local time), with the vast majority having temperatures over 26°C, and nowhere on the island had temperatures below 20°C.

By comparing the recorded temperature at Shannon Airport, which reached a maximum of 29.3°C to climatological data from the month of July we can see how abnormal these temperatures were. At all points during the day, the temperature was drastically higher than the 90th quantile for the month of July based on the previous thirty years; at it's lowest, it was still approximately 2.5°C warmer than the 90th percentile, a not insignificant number given that that occurred in the early morning, before the temperatures started rising with the increasing energy from the sun.

Even when compared to the 99th quantile, we can see that for most the day the July 22 temperatures were higher. There were a few points where the temperatures were equal to the 99th quantile, but both the maximum and minimum temperatures were markedly higher than the 99th quantile in general. It should also be noted that this data is slightly biased as our event is included in the data itself which is not ideal. If the entire 2021 heatwave

were to be excluded, the discrepancy between the historical data and the July 22 data would be even greater, further signifying how much of an anomalous event this heatwave was.

Looking at the forecasts, the first thing that is noticeable is that all of them under-predicted the temperature on average. The same-day, one-day and three-day forecasts were all close to each other, with their average under-prediction being within 0.5 °C of each other, but all were well outside the average temperatures. Looking at Figure 4, it can be seen that this happens because the forecasts predicted lower temperatures in the morning and a more rapid cooling as the day turned into night, neither of which actually occurred. By 00:00 on July 23, the forecasts were predicting a temperature under 17.5°C, whereas the actual temperature was still above $20°C$; the discrepancy in the forecasts was even greater in the morning, with even the best expecting the temperature would drop to 15°C, while the actual temperature hovered barely under 20°C at 05:00 UTC. However, despite the general under-prediction, the three closer forecasts accurately predicted the maximum temperature as well as the time it would occur at at.

Looking at both the data from the various forecasts at 15:00 UTC and their comparison at Shannon Airport, we can see that most of them did predict the correct temperature, to within a degree, at the hottest part of the day. Thus they were reasonably accurate to a first order. However, earlier in the day and as the day went on, the forecasts vastly under-predicted the temperature, with some discrepancies as high as 2.5°C. Even at 15:00, we can see that in some places the forecasts, except for the seven-day one, predicted the temperature would be warmer than it was, while other places it was predicted to be lower than it was.

The other thing to notice with the forecasts is that every single forecast model fared worse than just using a persistence-based forecast. The seven-day forecast, in particular, failed horribly at predicting the extreme temperatures Ireland would see, often tending closer to climatological averages rather than to the observed values or the other forecasts. This is likely because it was so far out, and the temperatures on the 14th of July, the last day before the forecast was started, were much cooler with no temperature at Shannon Airport reaching above 20°C. As more data came in, our forecasts got more accurate, with the latter three (three-day, one-day and same-day) tending to stay close together, especially the first two. While the same-day forecast tended to stay close to the persistence method forecast in most the skill scores, a sizable difference was noticed in the V score, with the per-

sistence model scoring 0.16 higher than the same-day forecast. This is not insignificant, and, coupled with the higher performance of the persistence model in the other skill scores, shows that a persistence model would have been preferable than any of the forecasts in this instance.

One interesting peculiarity that does stand out is that it seems that the temperature forecast is one hour ahead of the actual temperatures. If we shift the observed temperatures back an hour, a plot like Figure 5 is produced, with significantly more overlap during the hours from 10:00 to 20:00. Comparing the skill scores with this shift shows an increase in accuracy in all of the latter three forecasts though they still under-perform in the morning. During the period from 10:00 to 20:00 UTC the ACC score of the forecast increases to 0.998, while the average bias during those 10 hours drops to -0.3, a significant drop. There are several possible reasons this error could have been caused. One is due to a possible indexing error somewhere in the Shannon Airport data not compensating for Irish Summer Time (UTC+1). If, as unlikely as it is, IST readings of the time had been used, this would have resulted in the shift we saw. Another possibility is the difference in how quickly the air mixes in the WRF model versus how it mixes in real life. As the sun heats the ground which then heats the air, this could speed up or slow down the speed at which the forecast predicted certain temperatures would be reached, as well as how quickly the air would cool at the end.t

Figure 5: July 22, 2021 (shifted backwards by one hour) compared to forecasts at Shannon Airport

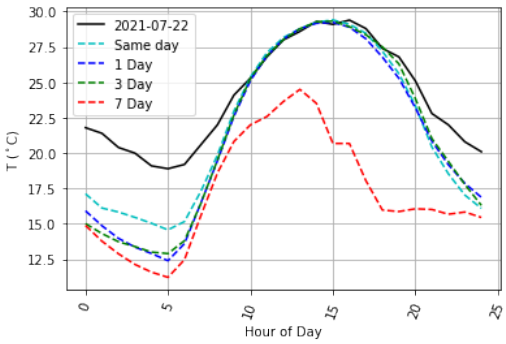

Overall, we can see that the heatwave which peaked on July 22, 2021 was an extremely abnormal event. It was far outside traditional climatological values, falling above the 99th percentile in regular climatological data of

temperatures. Likewise, all forecasts models, including that one run starting on the same-day and the day before failed to predict the duration of the heat, having it drop off much quicker than it actually did. This thus caused all their skill scores to be lower than that from persistence modelling. Due to the anomalous nature of this event, the best indicator of the temperatures that were reached during it continues to be persistence forecasting.