Automatic Gravity Wave Detection & Characterisation from NWP model output

Jonathan Coney[1], Andrew Ross[1], Leif Denby[1], He Wang[2], Simon Vosper[3], Annelize van Niekerk[3], Tom Dunstan[3]

[1]ICAS, School of Earth and Environment, University of Leeds,
[2]School of Computing, University of Leeds,
[3]Met Office, Exeter

ICAS Internal Seminar 25 April 2022

Mountain Waves

UNIVERSITY OF LEEDS

- Vertically Propagating Waves & Trapped Lee Waves
- Both gravity waves, in this case formed by the forced ascent of air over orography.

Teixeira (2014)

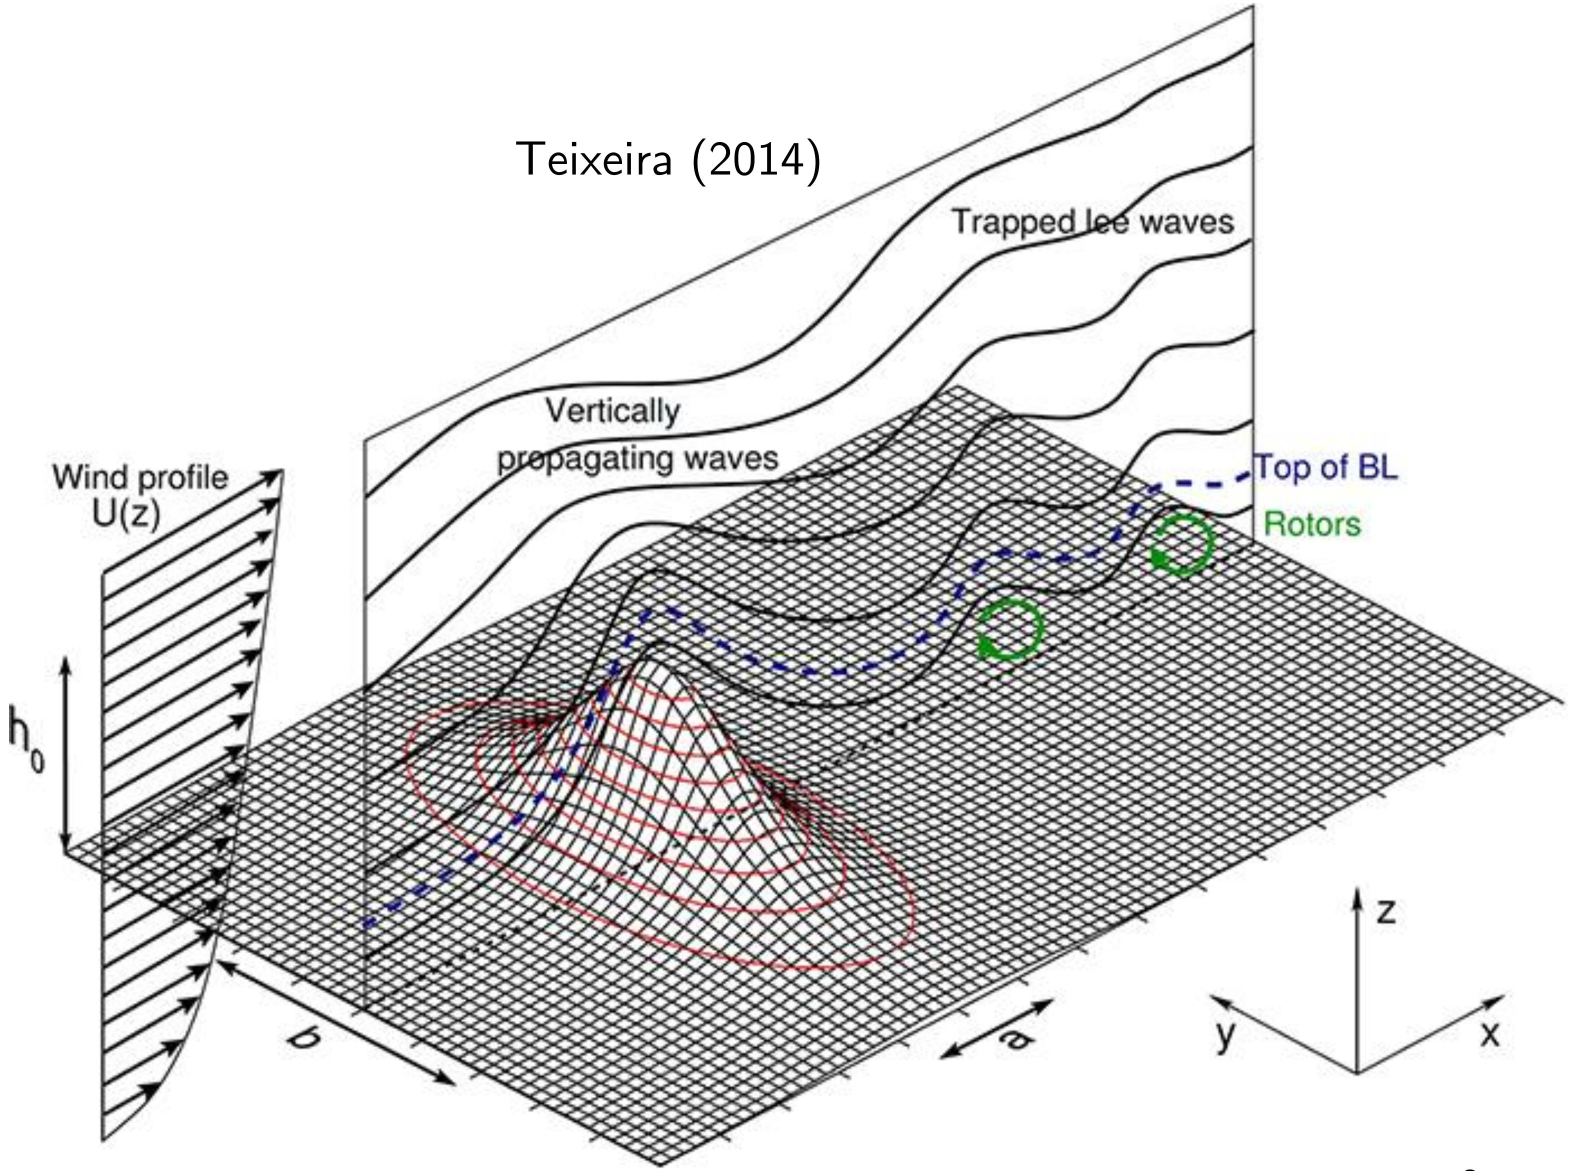

Forecasting mountain lee waves

- Why do we care?
  - Strong downslope winds (ie. Sheffield windstorm 1962)
  - Turbulent rotors (e.g. BOAC Flight 911)

Yorkshire Post

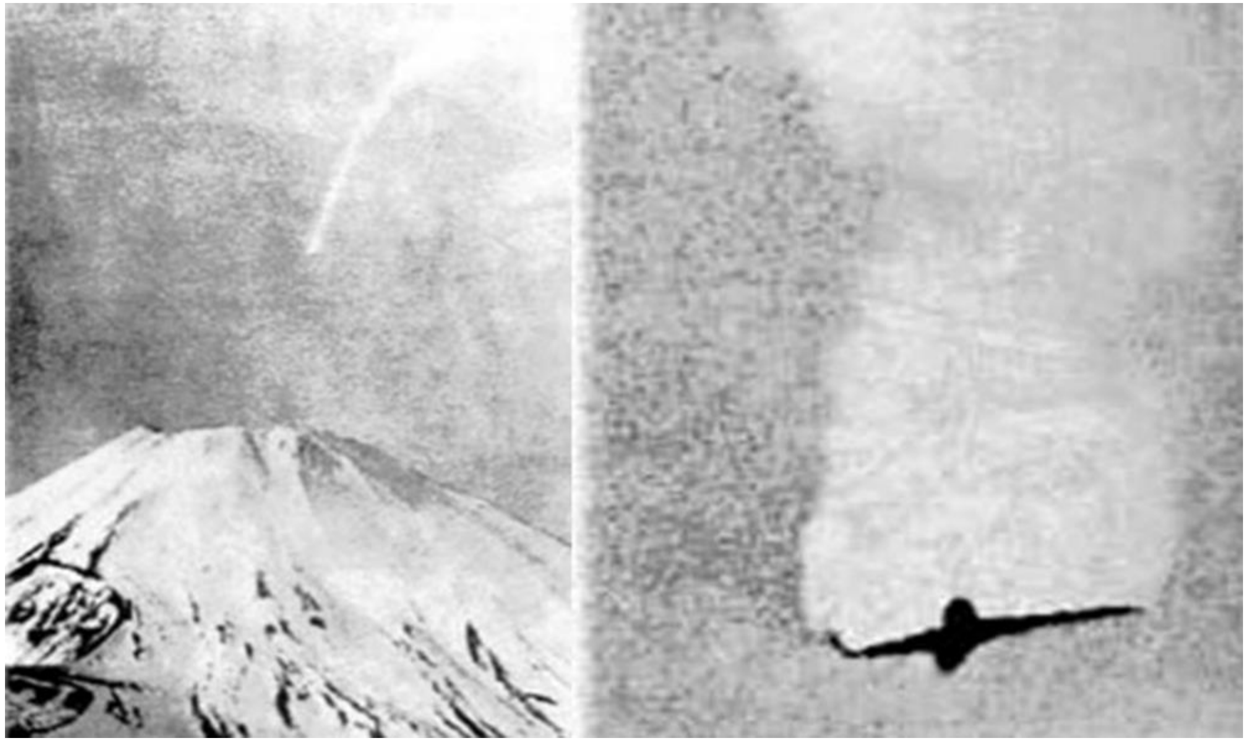

Bureau of Aircraft Accidents Archives

UKV developments

- UKV is the Met Office's high resolution forecast model for over the UK
- Horizontal spatial resolution ~2km.
- Since upgrade to dynamical core: Even Newer Dynamics for General atmospheric modelling of the environment (ENDGame) - UKV now resolves lee waves.

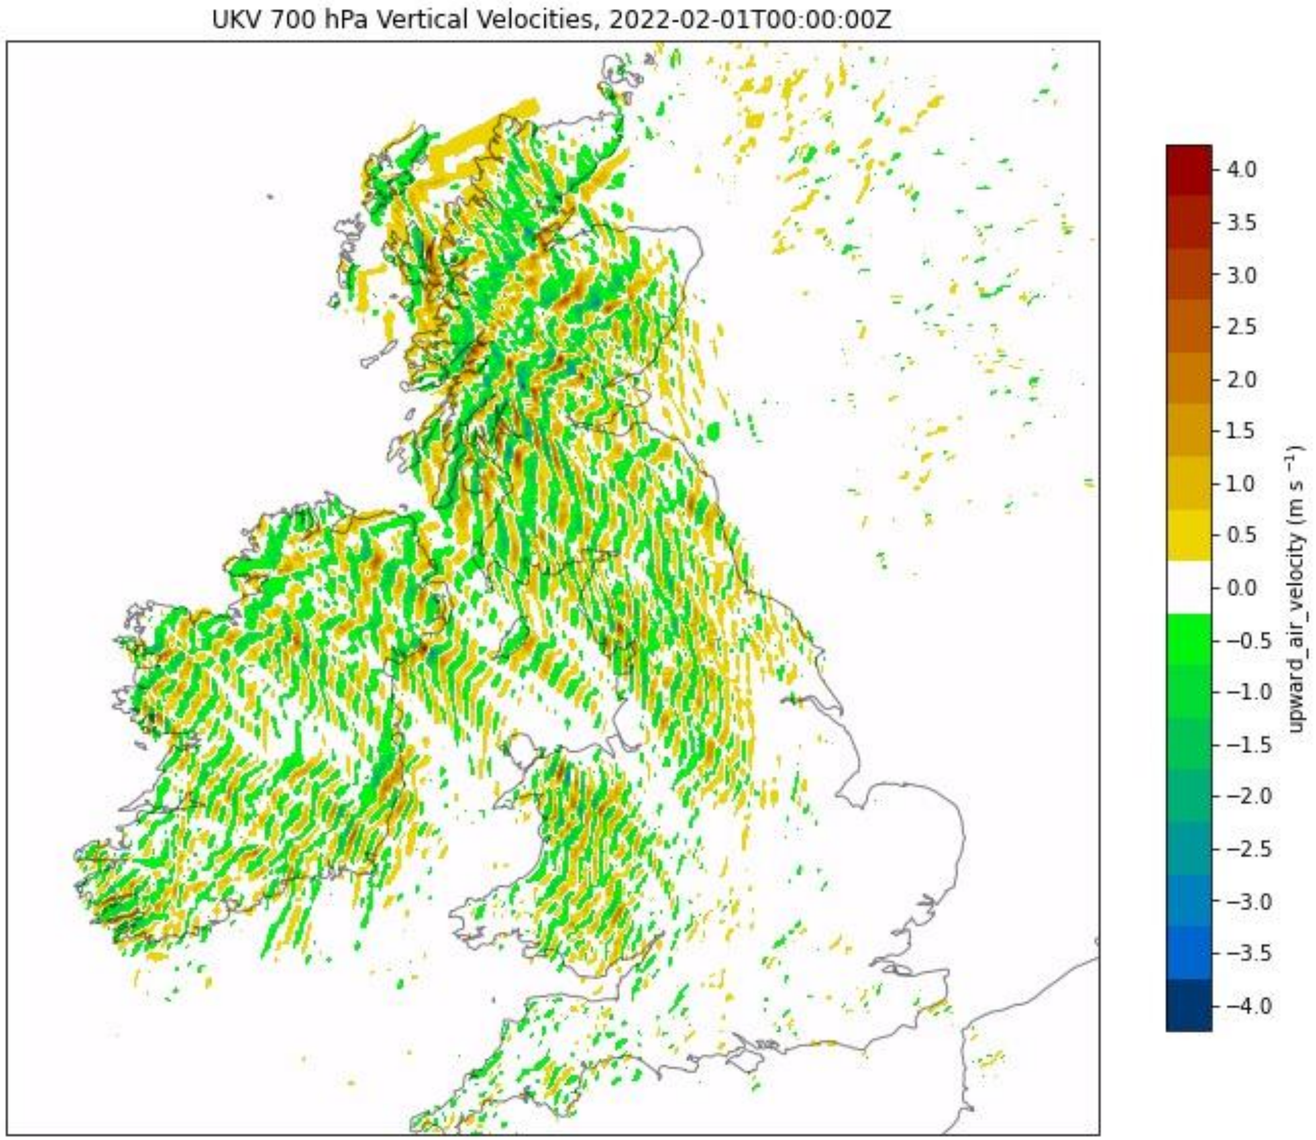

Lee Waves in UKV model output at 700 hPa.

Method: Lee Wave Recognition

UNIVERSITY OF LEEDS

- How can we use that lee waves are resolved in UKV to learn more about them? First we need a way to recognise them in model output: **machine learning.**
- Deep Learning using fastai.
- U-Net to pixelwise-classify (segment) some UKV output into regions containing waves and those that don't.
- **Supervised learning**, so we need some labelled data to create training data.
- Data is augmented during training in an attempt to train a more generalisable model.

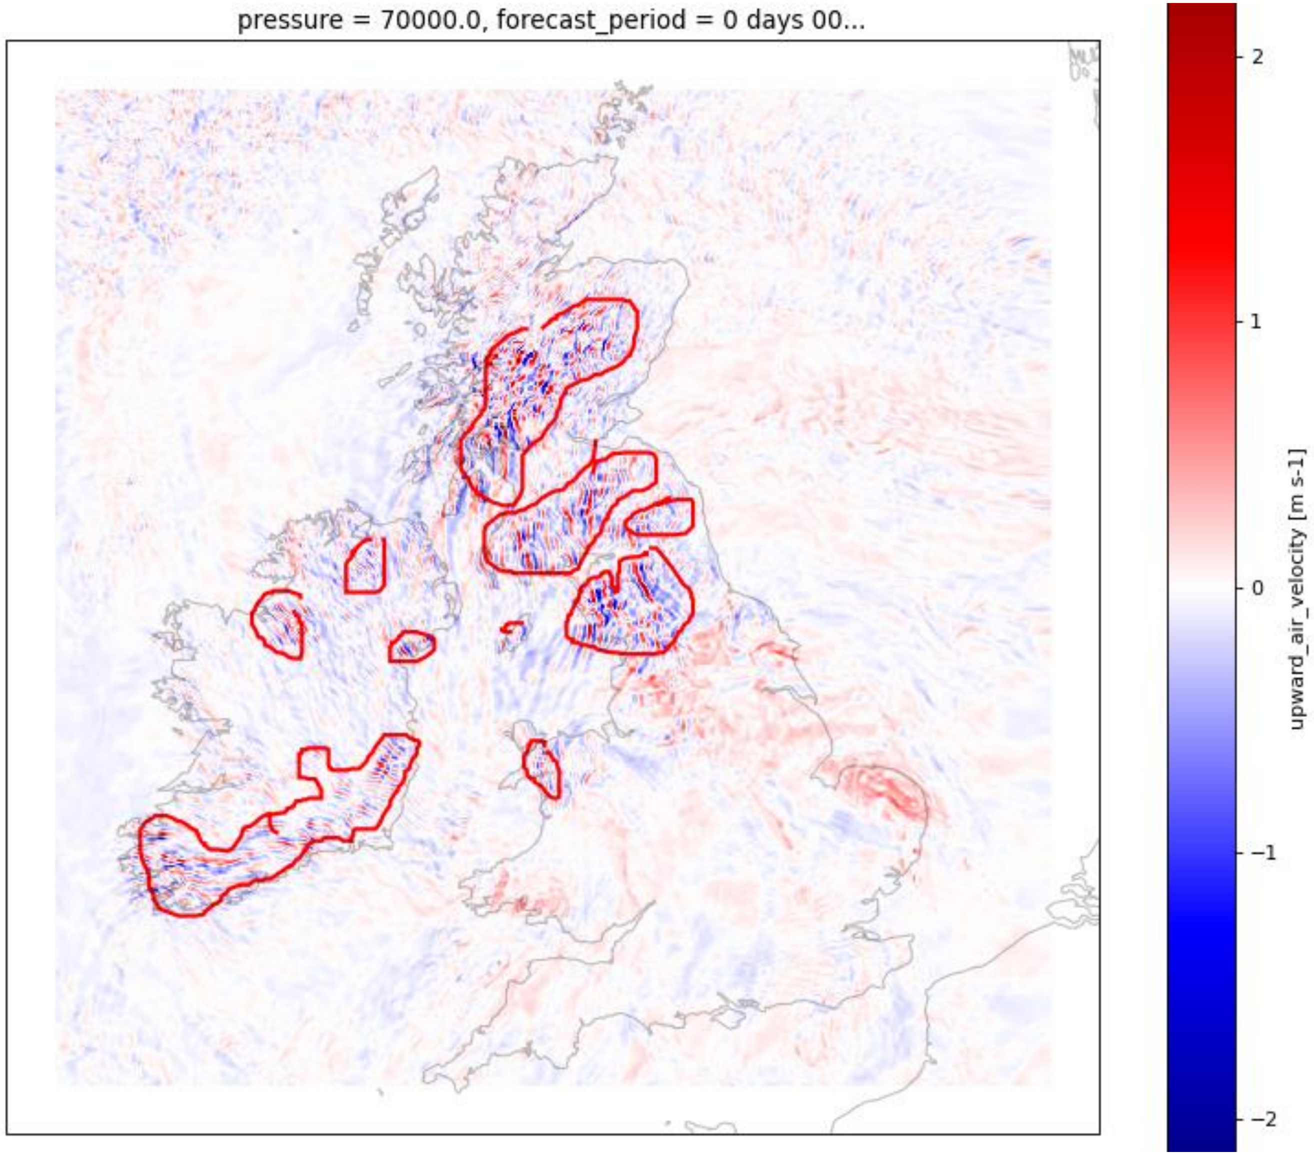

Method: Lee Wave Recognition

- How can we use that lee waves are resolved in UKV to learn more about them? First we need a way to recognise them in model output: **machine learning.**
- Deep Learning using fastai.
- U-Net to pixelwise-classify (segment) some UKV output into regions containing waves and those that don't.
- **Supervised learning**, so we need some labelled data to create training data.
- Data is augmented during training in an attempt to train a more generalisable model.

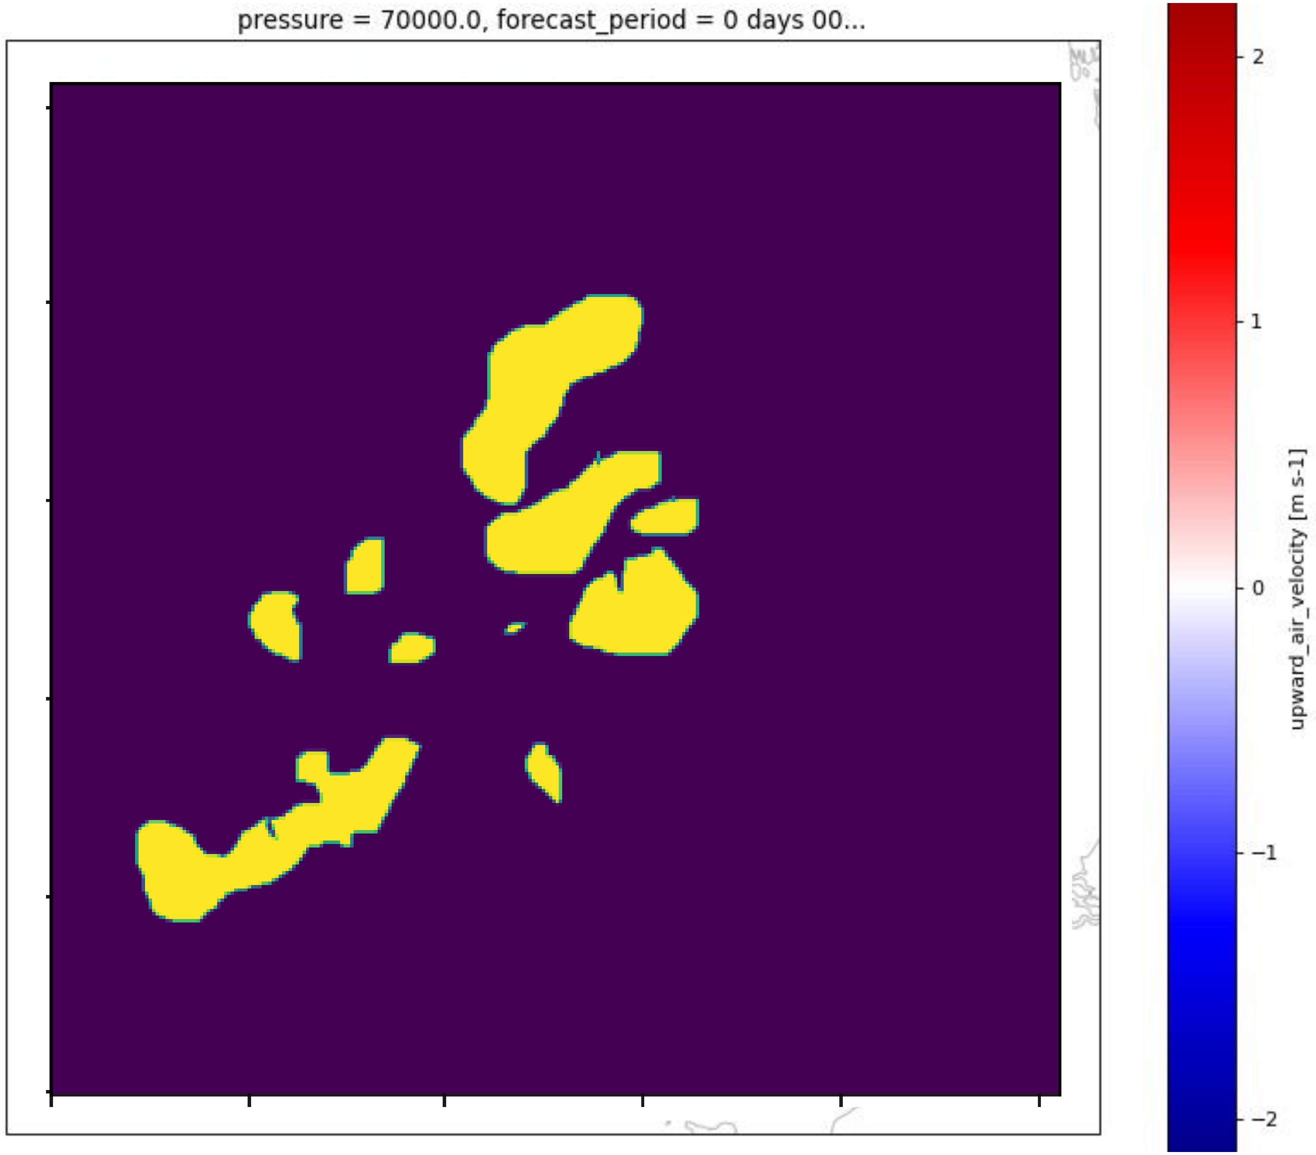

The ML model: A U-Net

- Output has the same shape as the input
- Originally used for biomedical images

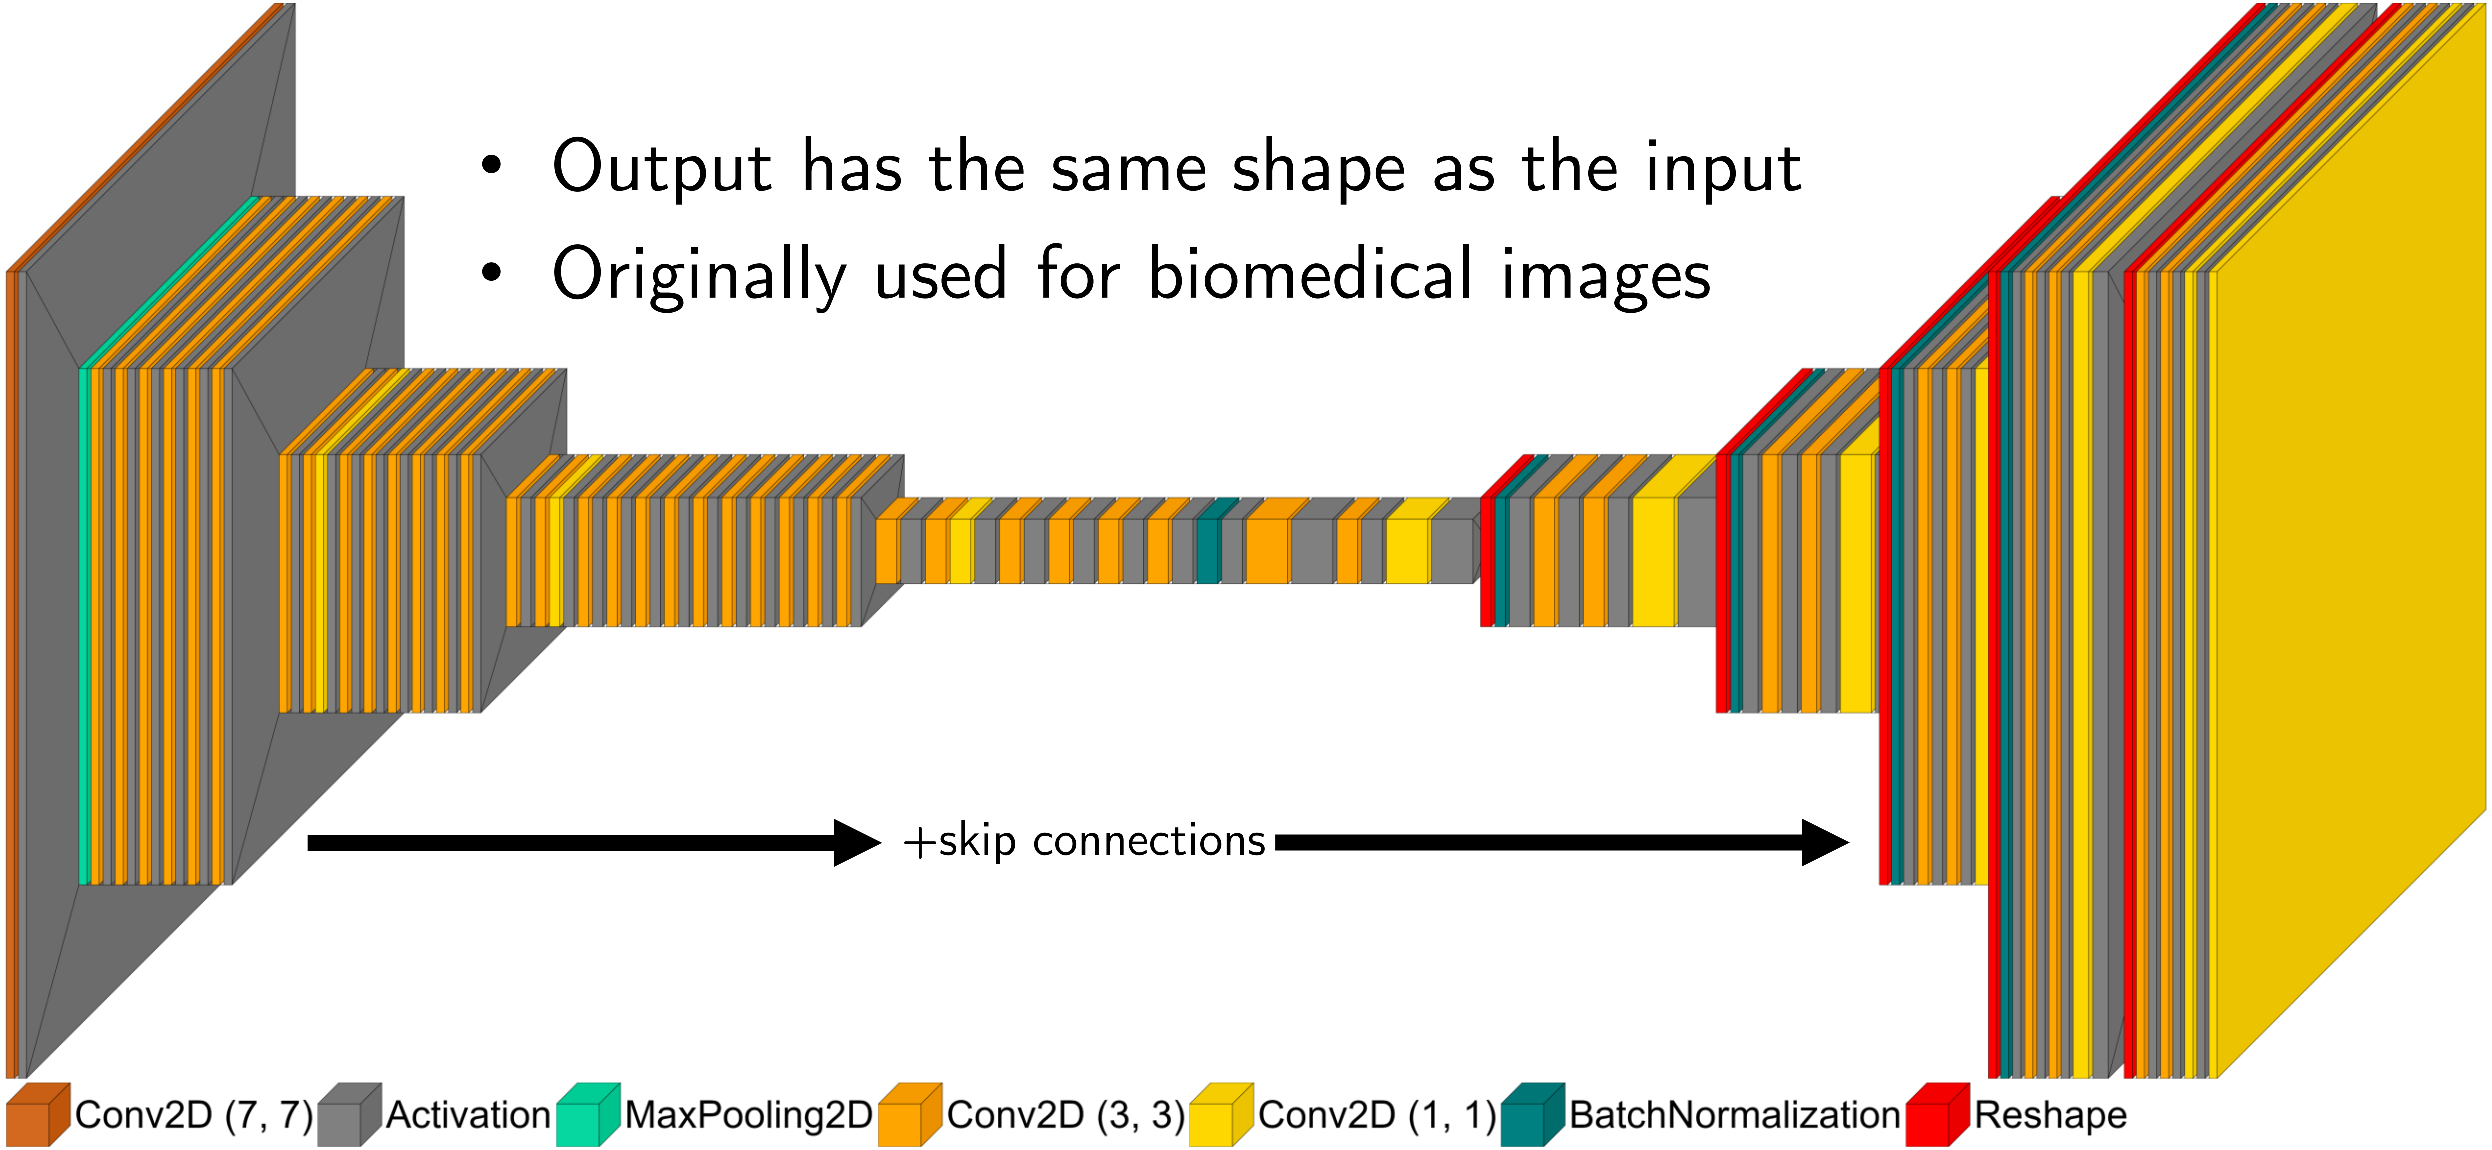

Some example results on the test set

- On average on the test set: 91% pixel accuracy; 0.68 IoU score

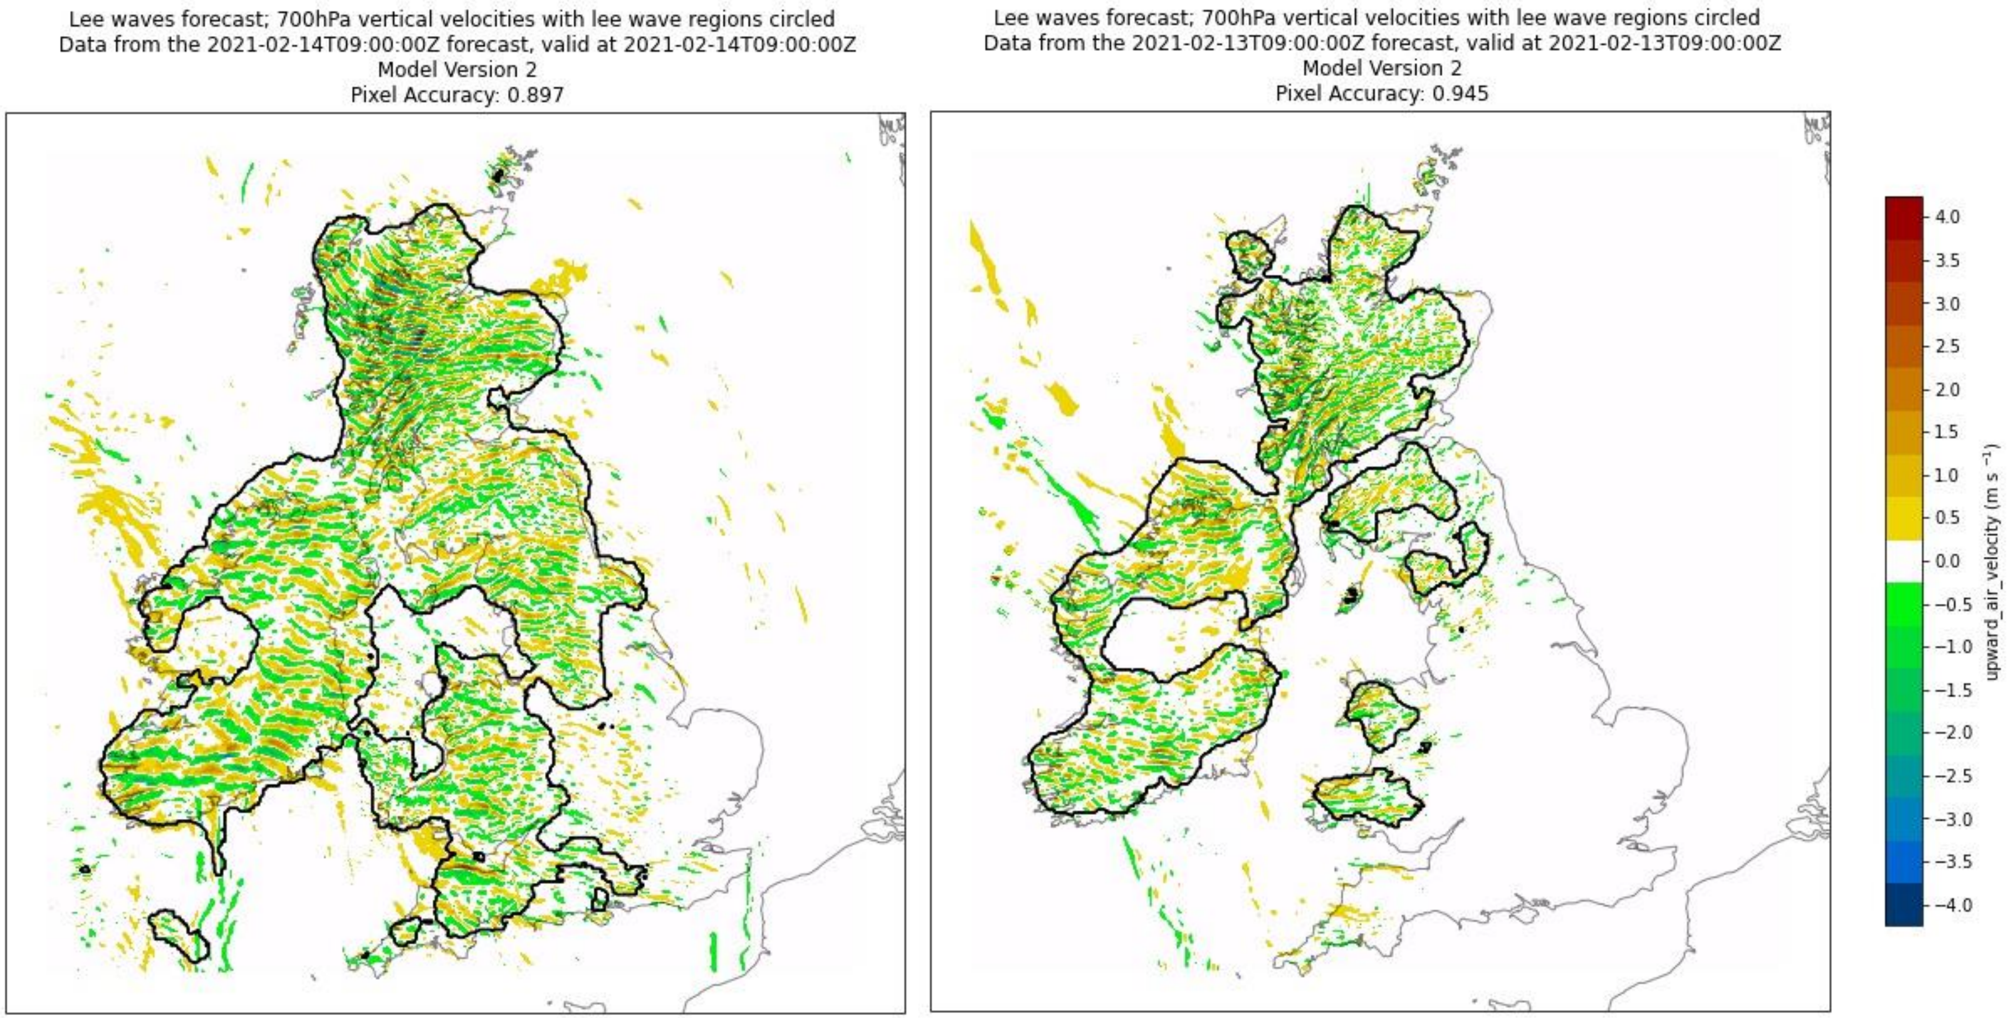

IoU Score = Intersection of "truth" and prediction ÷ Union of "truth" and prediction.

Want to get the IoU as close to 1 as possible.

Crowdsourcing Test Data labels using Zooniverse

- Evaluation of model predictions versus hand-labelled truth.
- Someone generously suggested that my model might be better than my hand-labels.
- So crowdsource the labelling process, with volunteers here and at the Met Office
- Need >60% of labellers to agree for a pixel to be a wave.

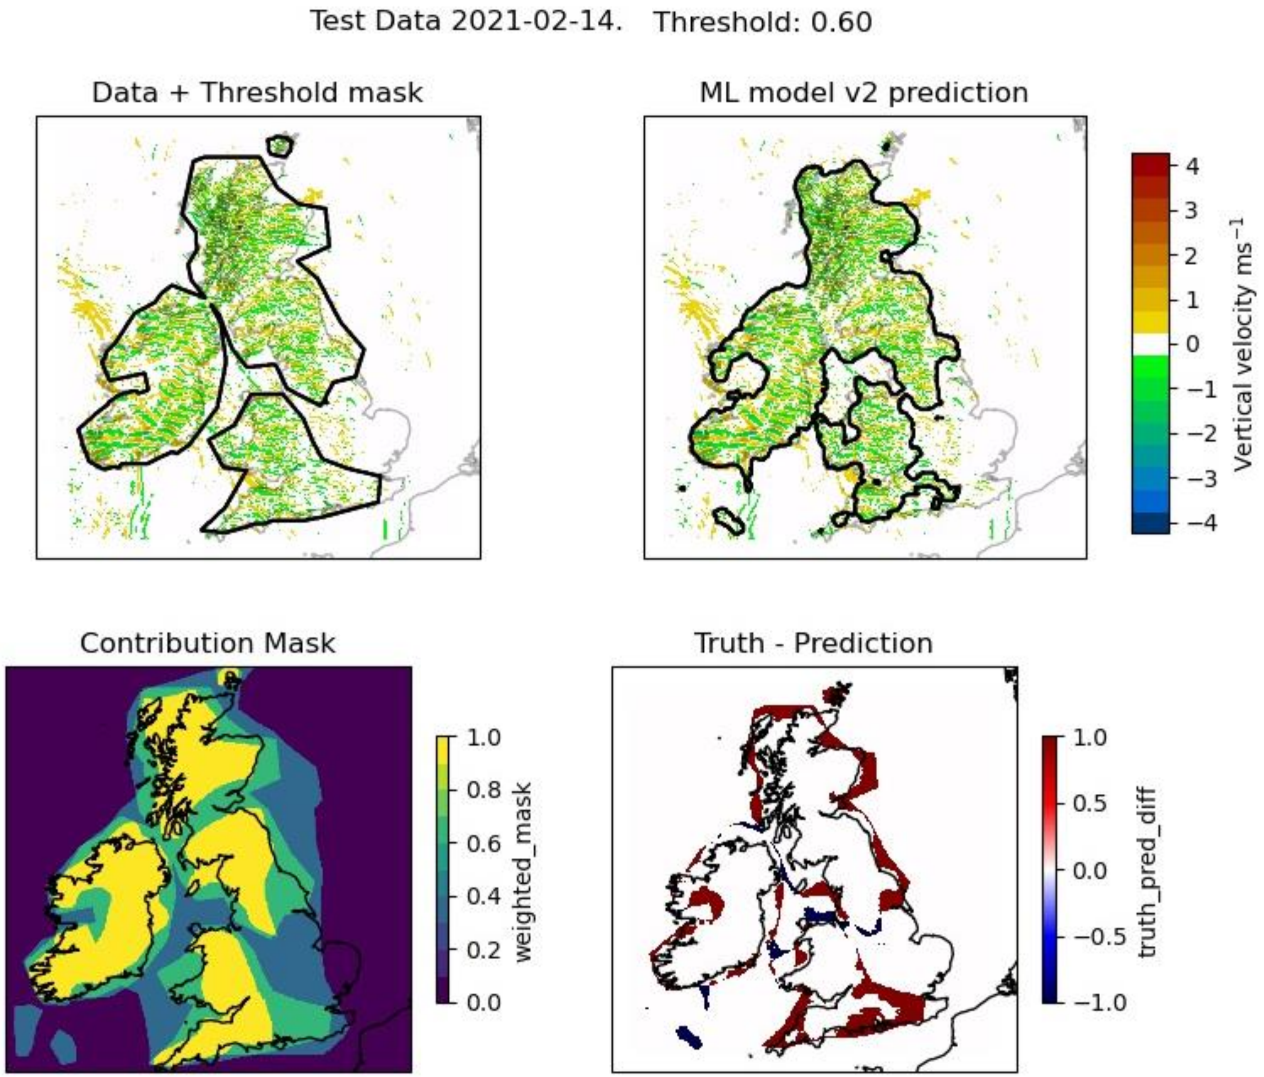

Zooniverse labels on the test set. 60% agreement

UNIVERSITY OF LEEDS

Just my "truth"

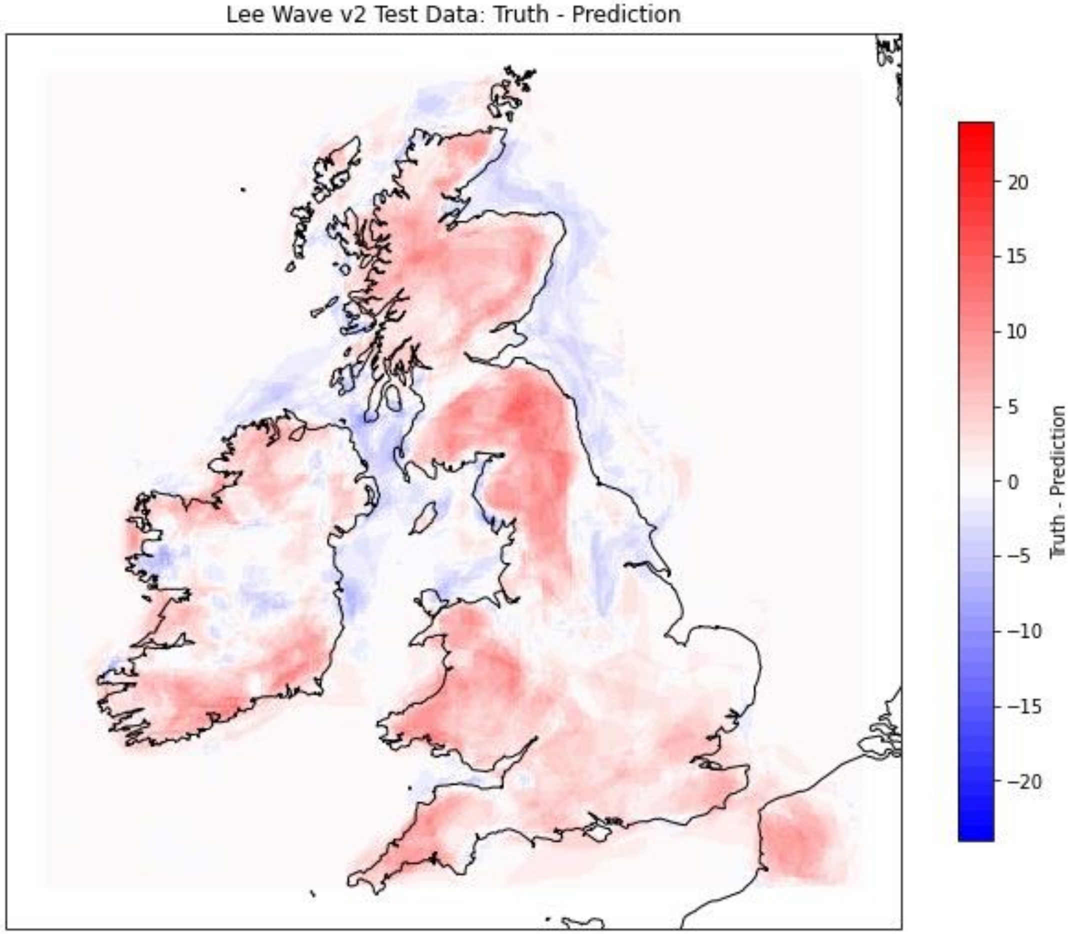

Combined truth using Zooniverse labels

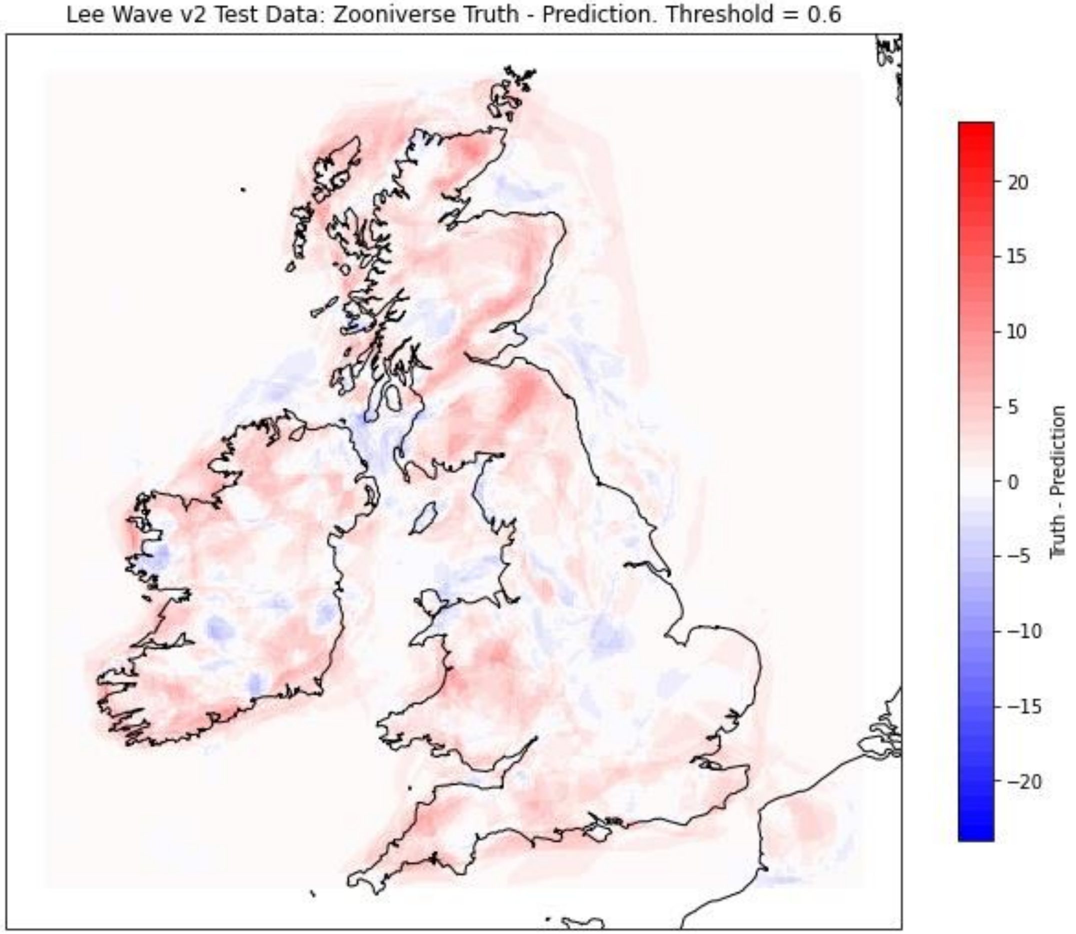

Compared to "truth": Red: Model underpredicts waves. Blue: Model overpredicts waves

What has the ML model actually learned?

UNIVERSITY OF LEEDS

Use GradCAM to analyse the model's predictions. The model has learned to penalise what it recognises as convection.

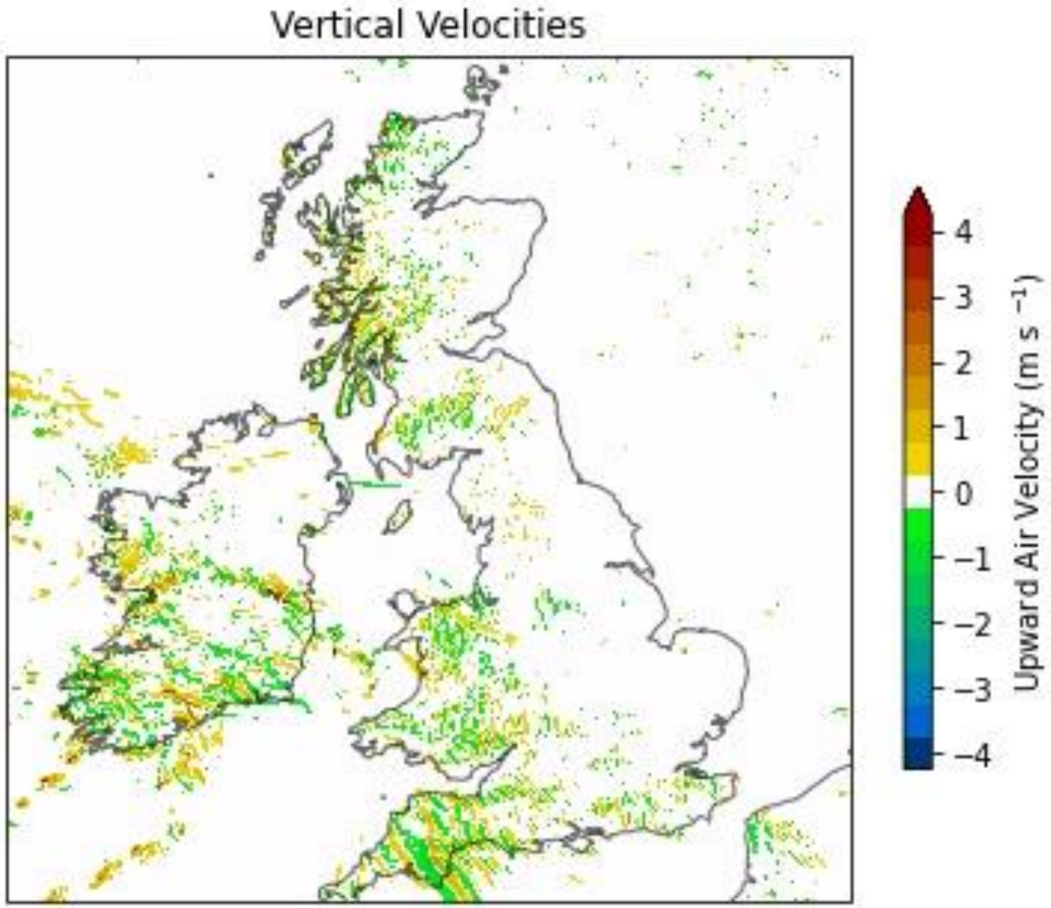

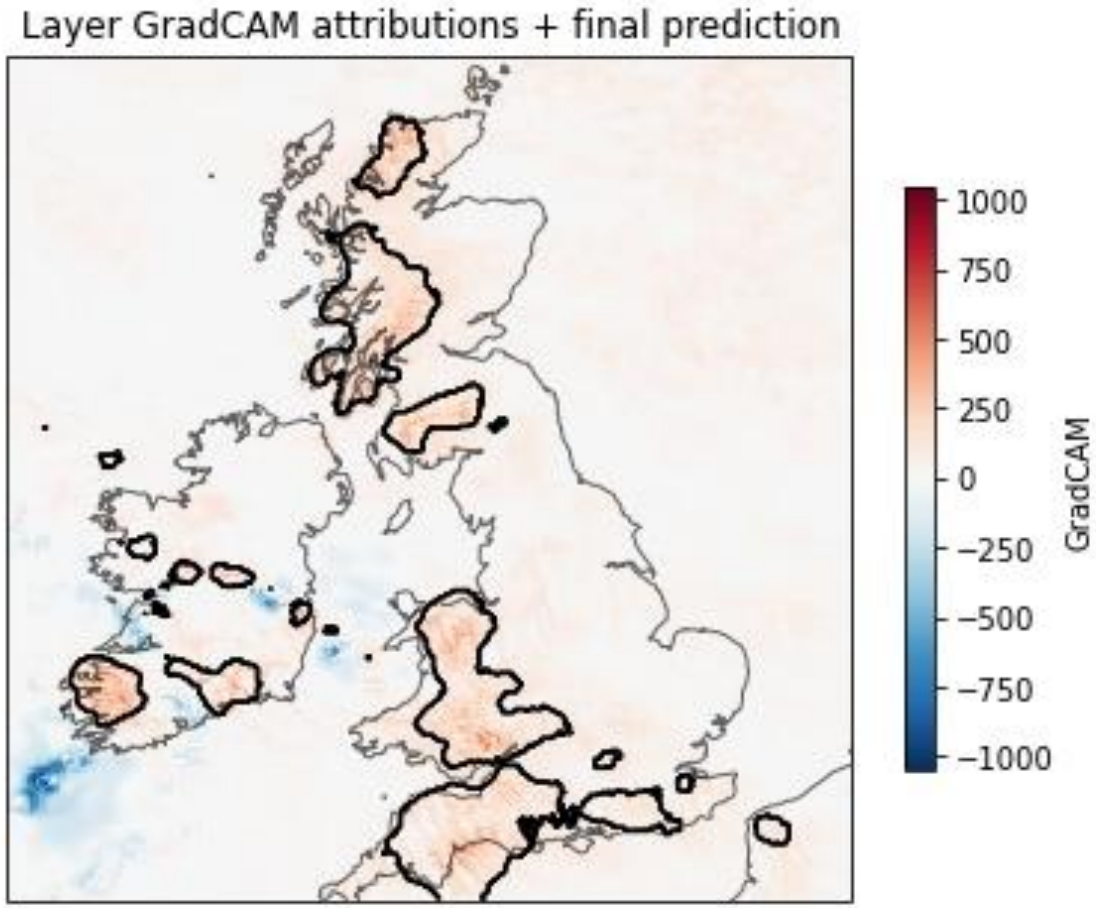

What has the ML model actually learned?

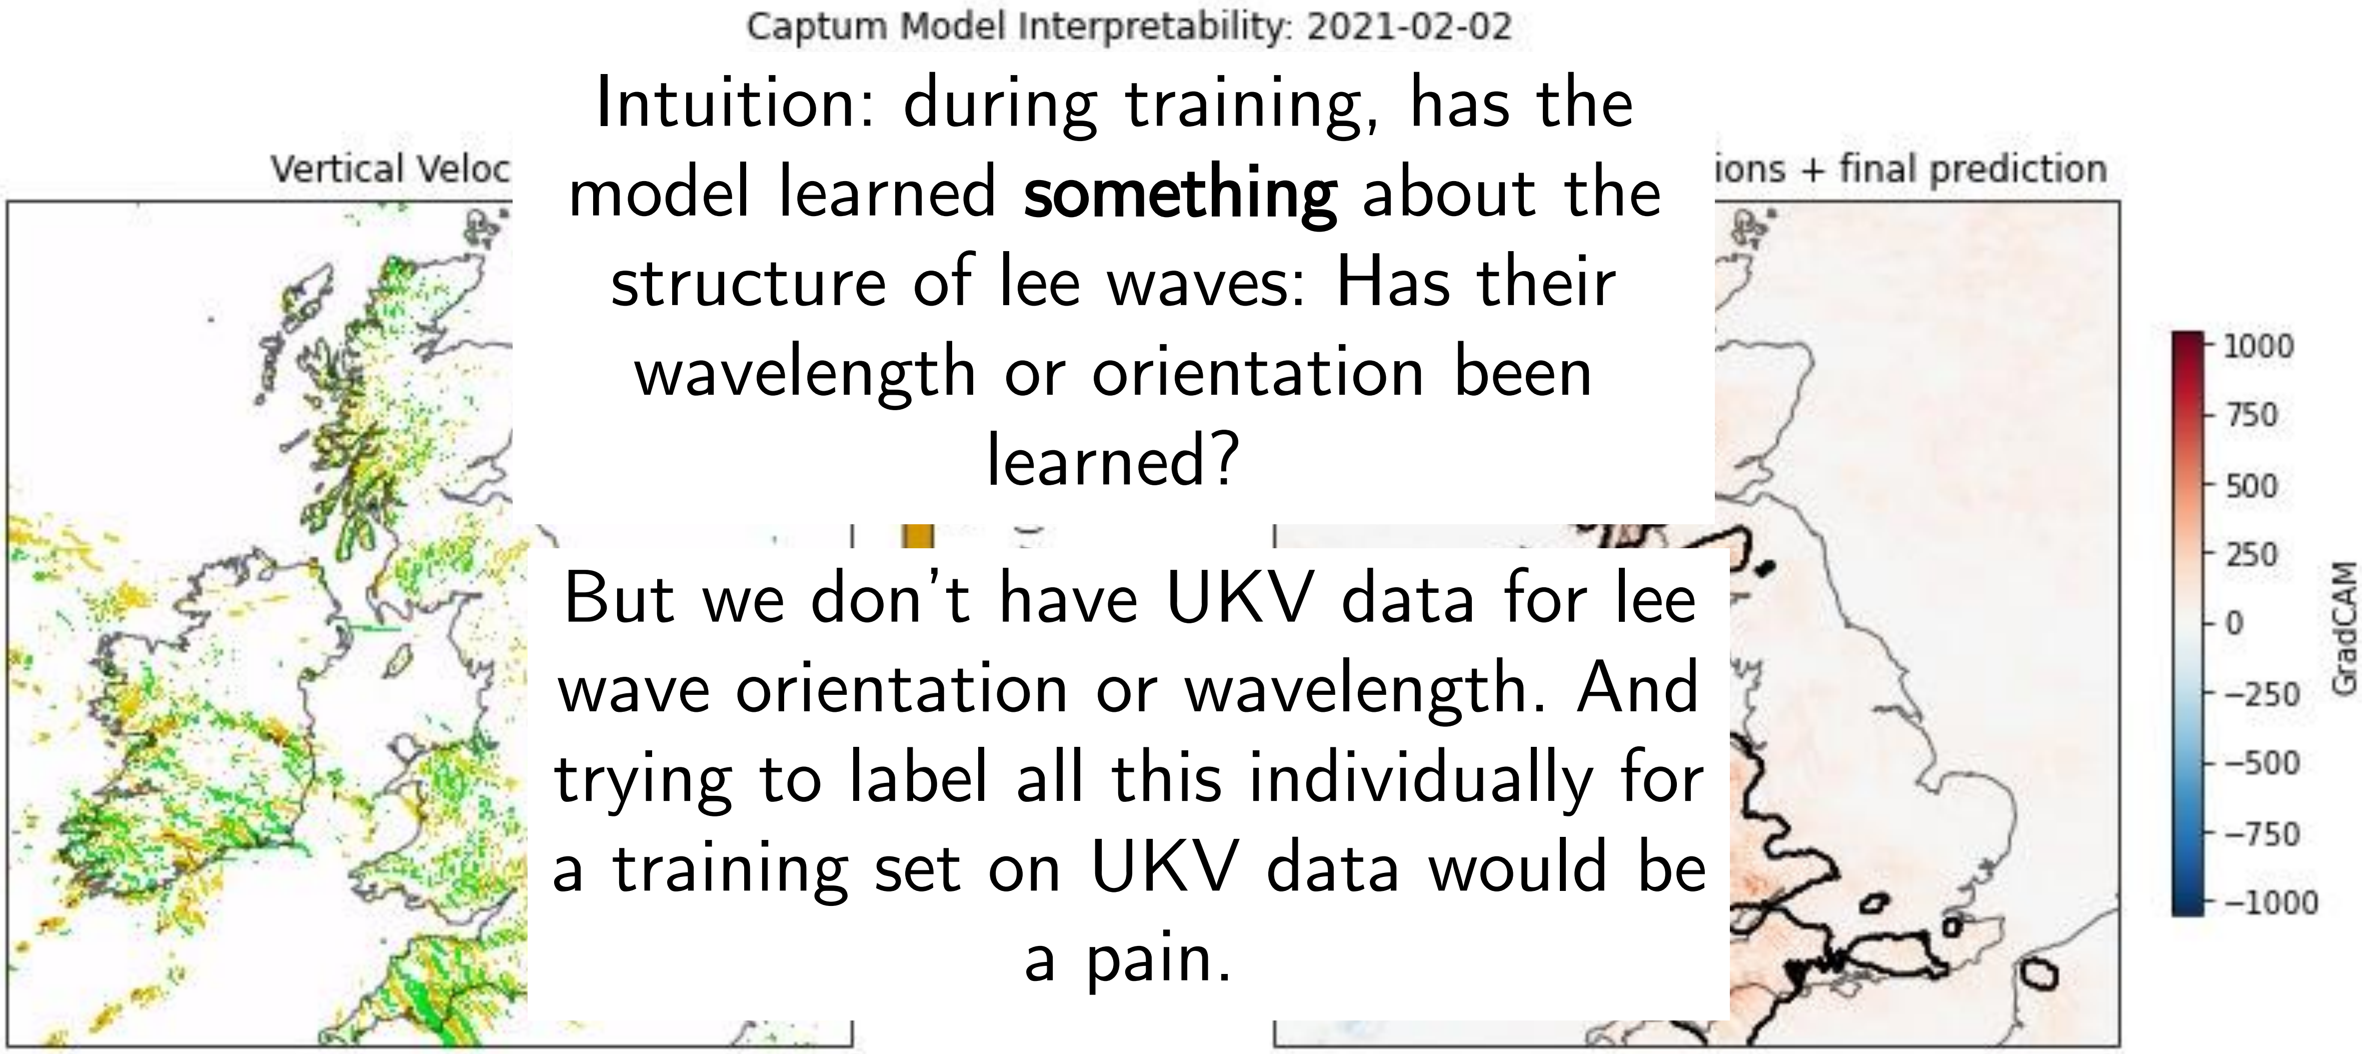

Intuition: during training, has the model learned **something** about the structure of lee waves: Has their wavelength or orientation been learned?

But we don't have UKV data for lee wave orientation or wavelength. And trying to label all this individually for a training set on UKV data would be a pain.

Solution: construct some synthetic data

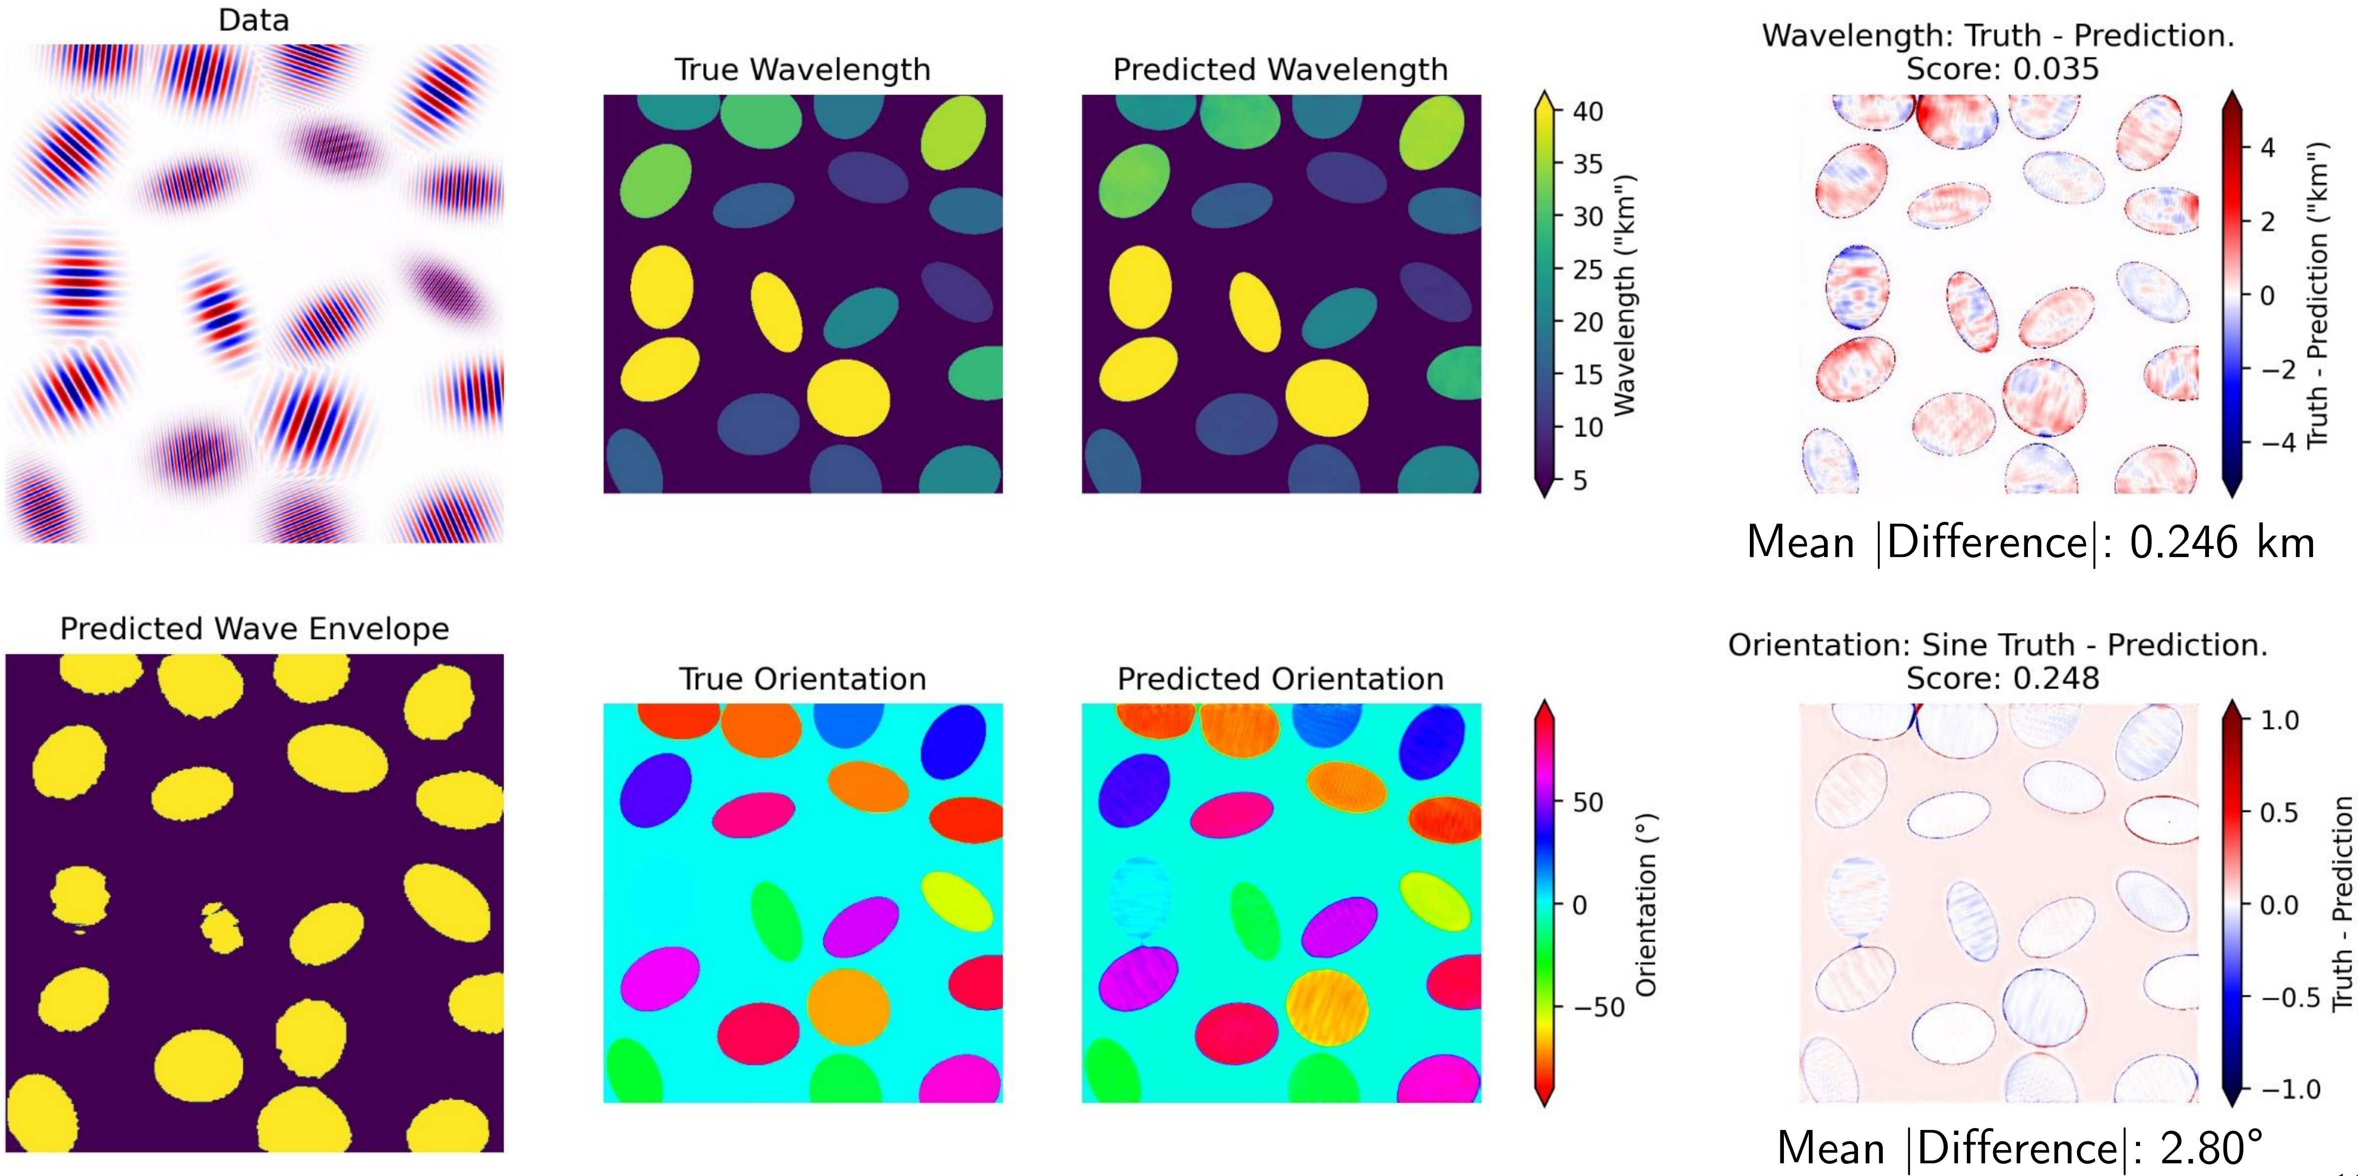

U-Net summary – transfer learning

- Freeze the majority of the model
- Only finetune the last few layers (dashed box)

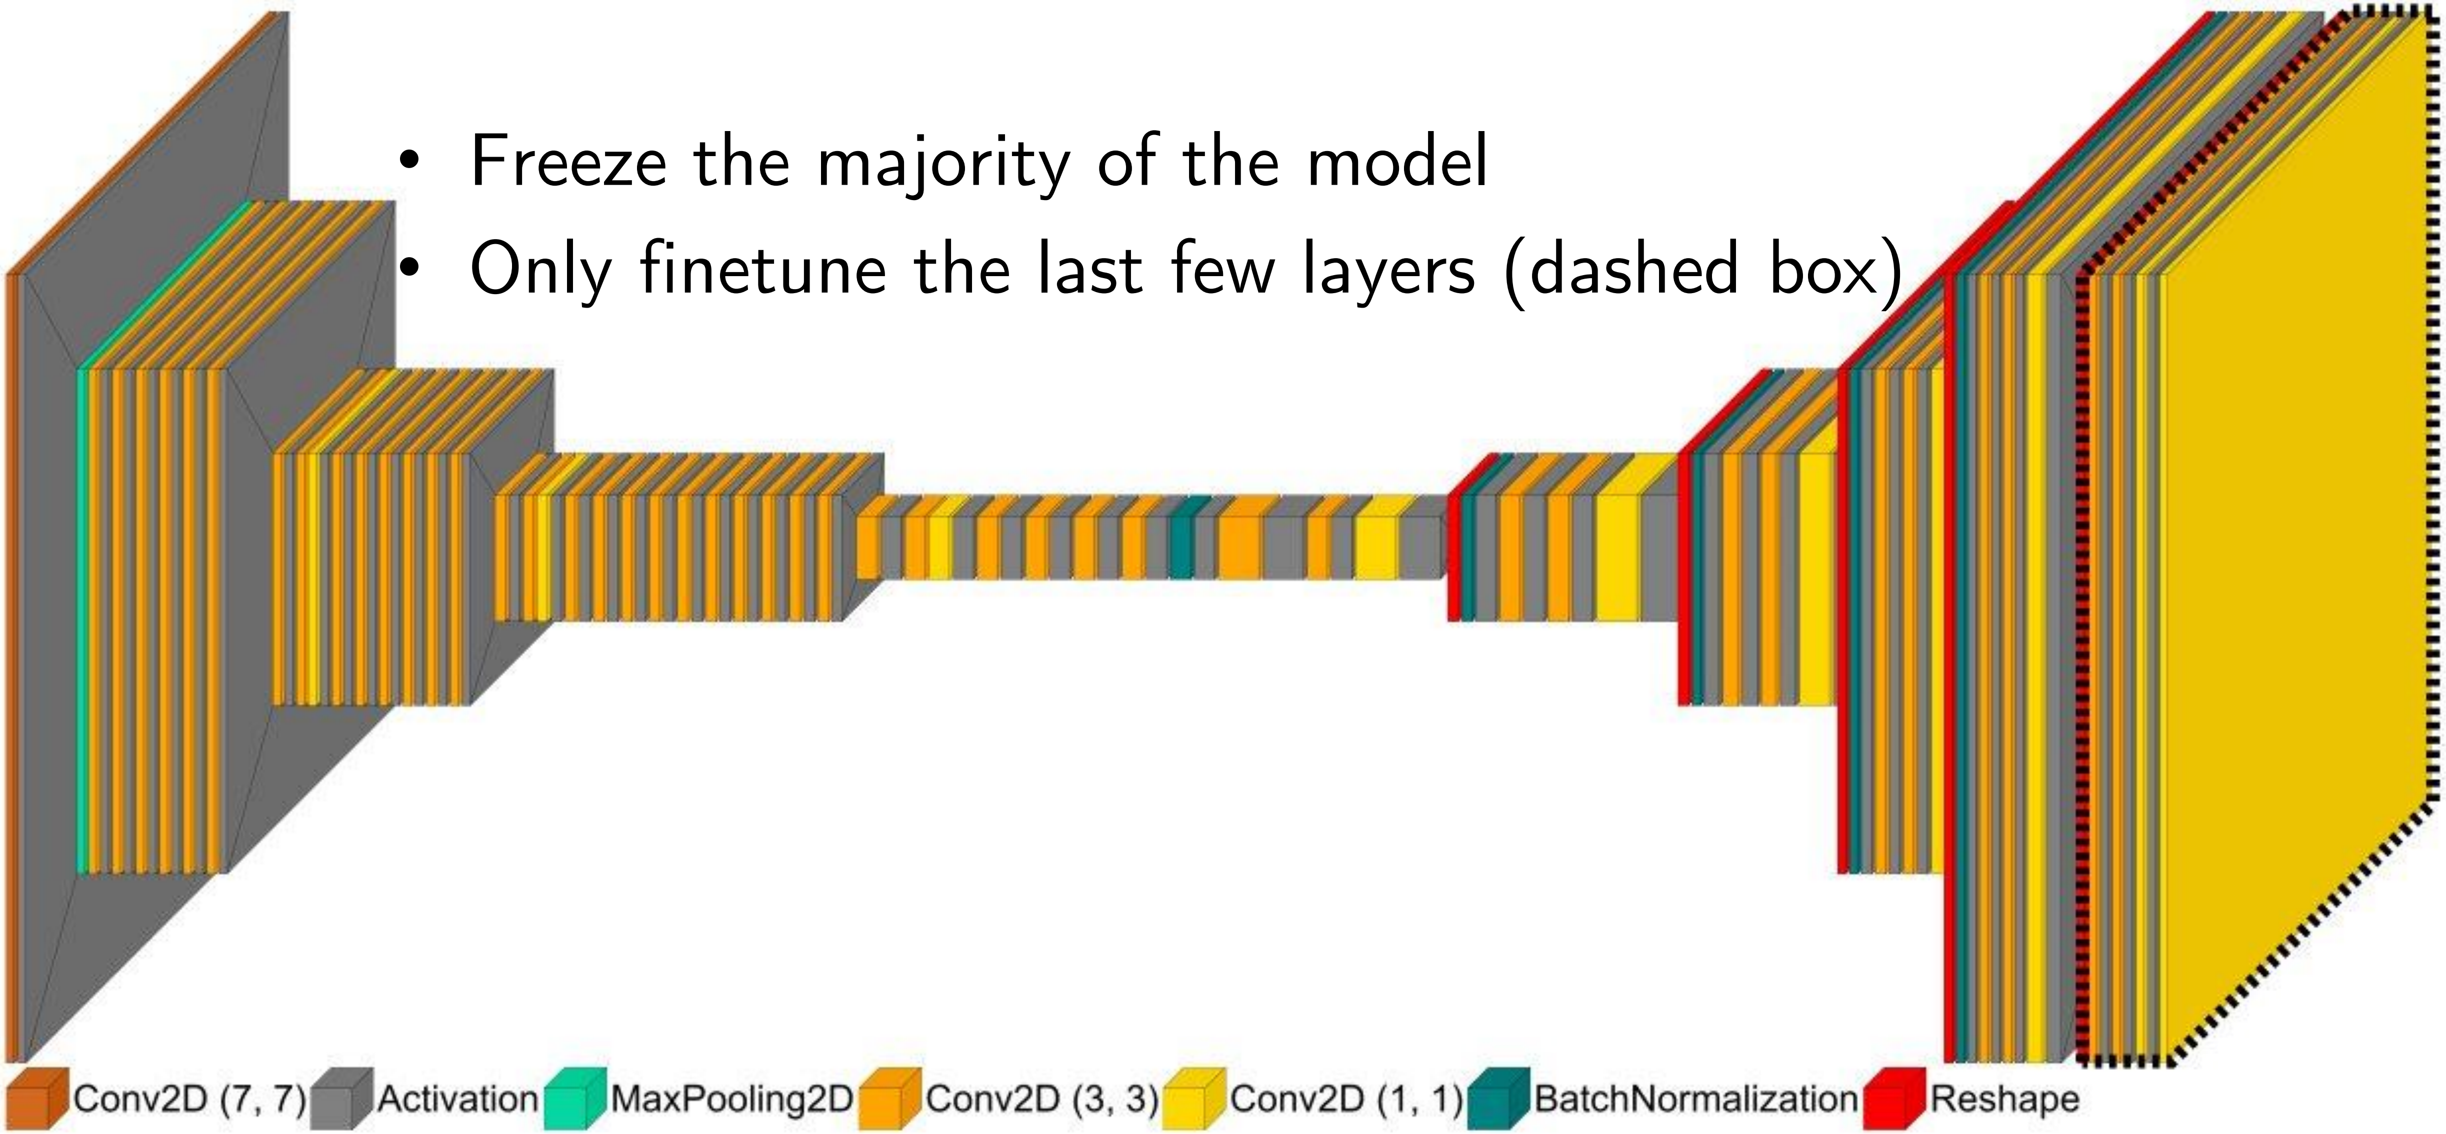

Performance on the UKV data

UNIVERSITY OF LEEDS

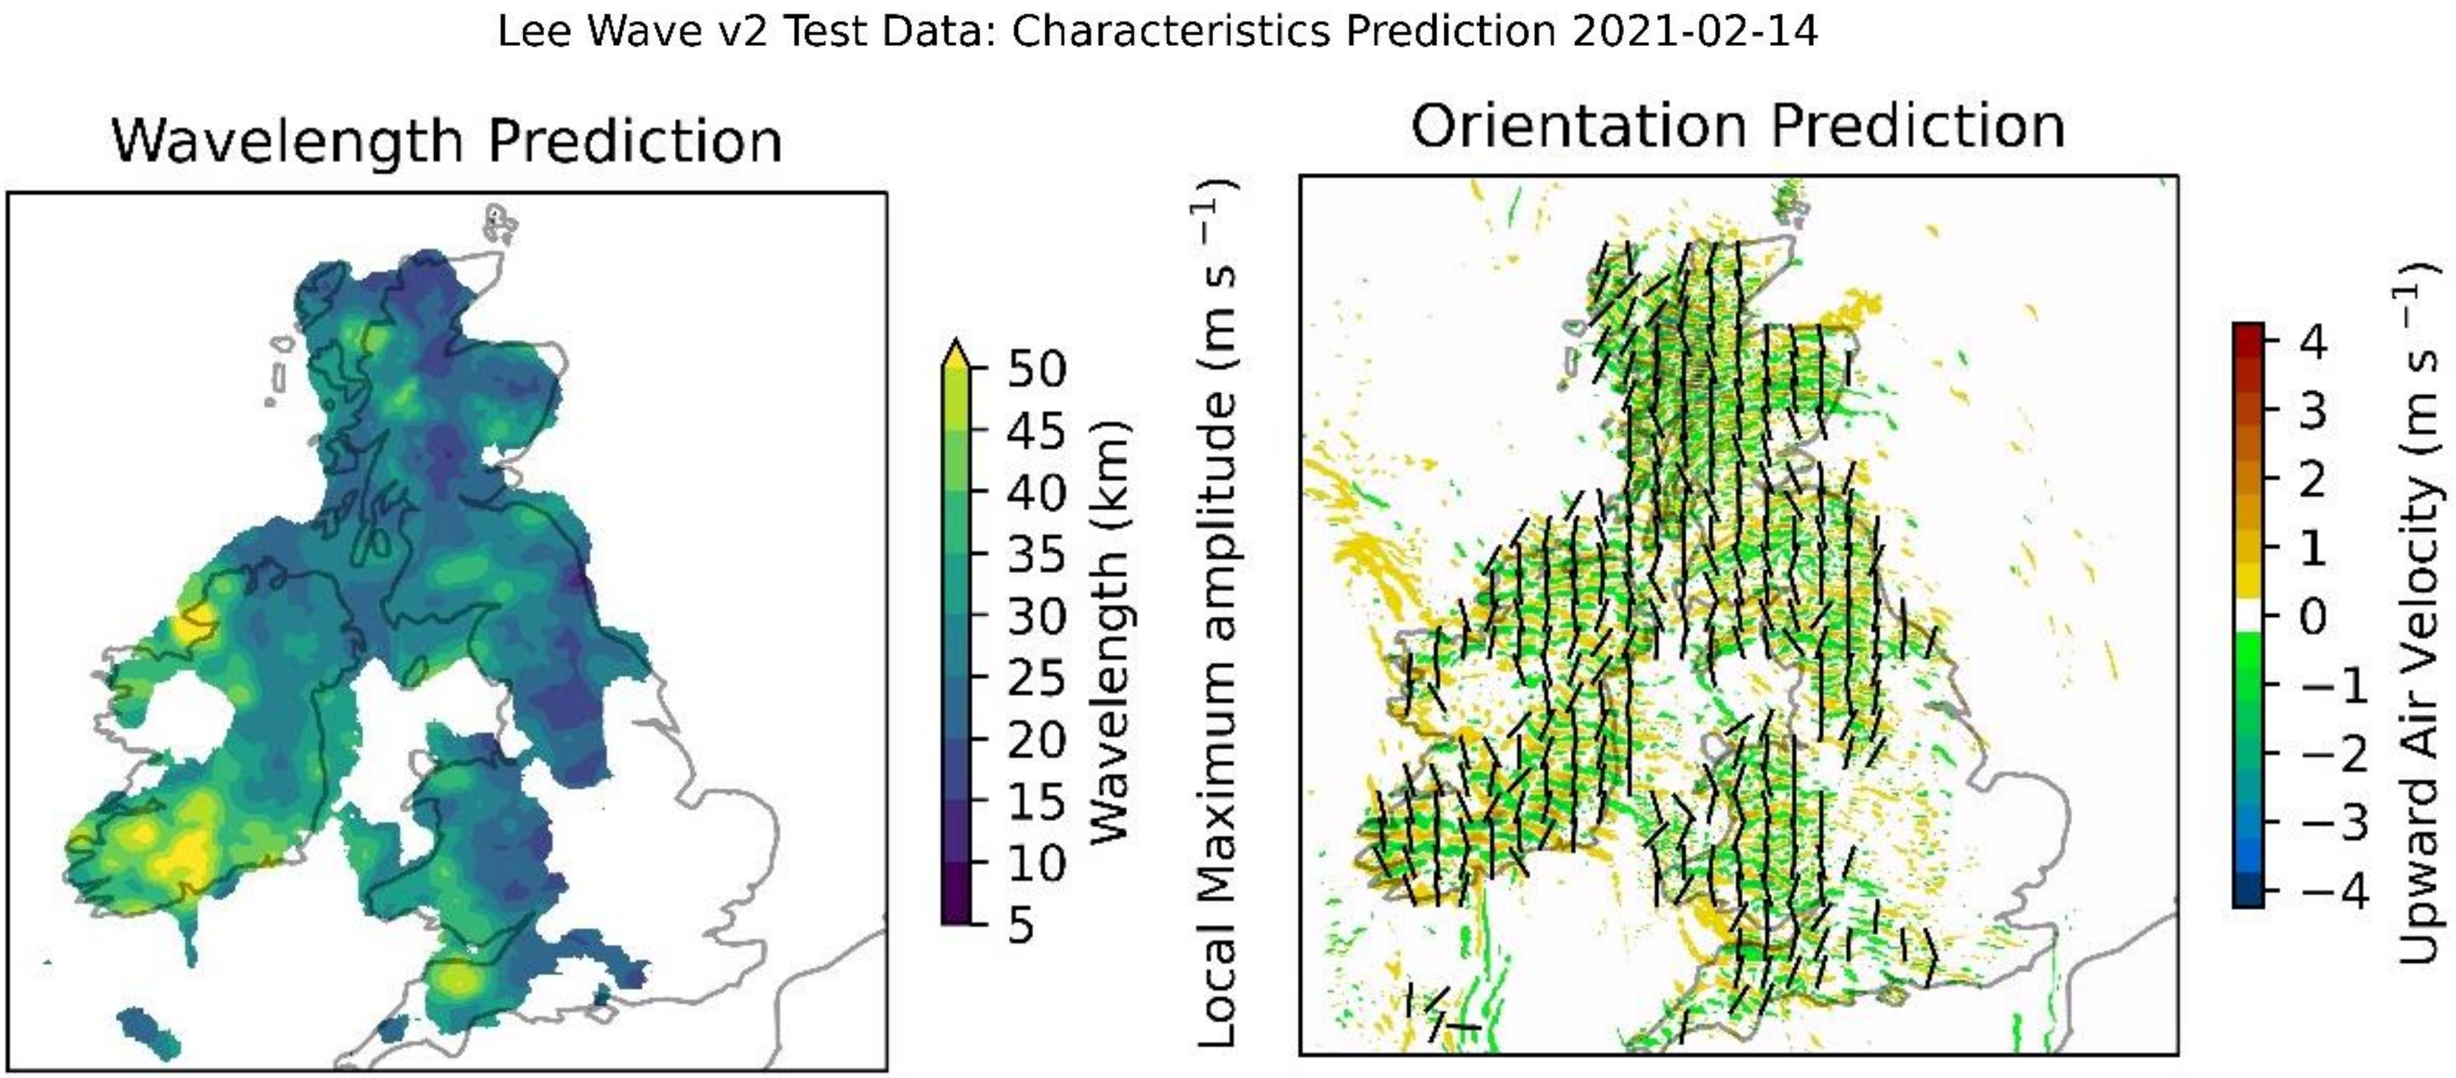

But how do we verify results on the actual UKV data?

- Wavelength – a wavelet analysis of some description? Still to nail down.
- Orientation: Can compare against UKV Wind Directions (which is a decent approximation, but not perfect): 69% of waves within 30 ° in test set.

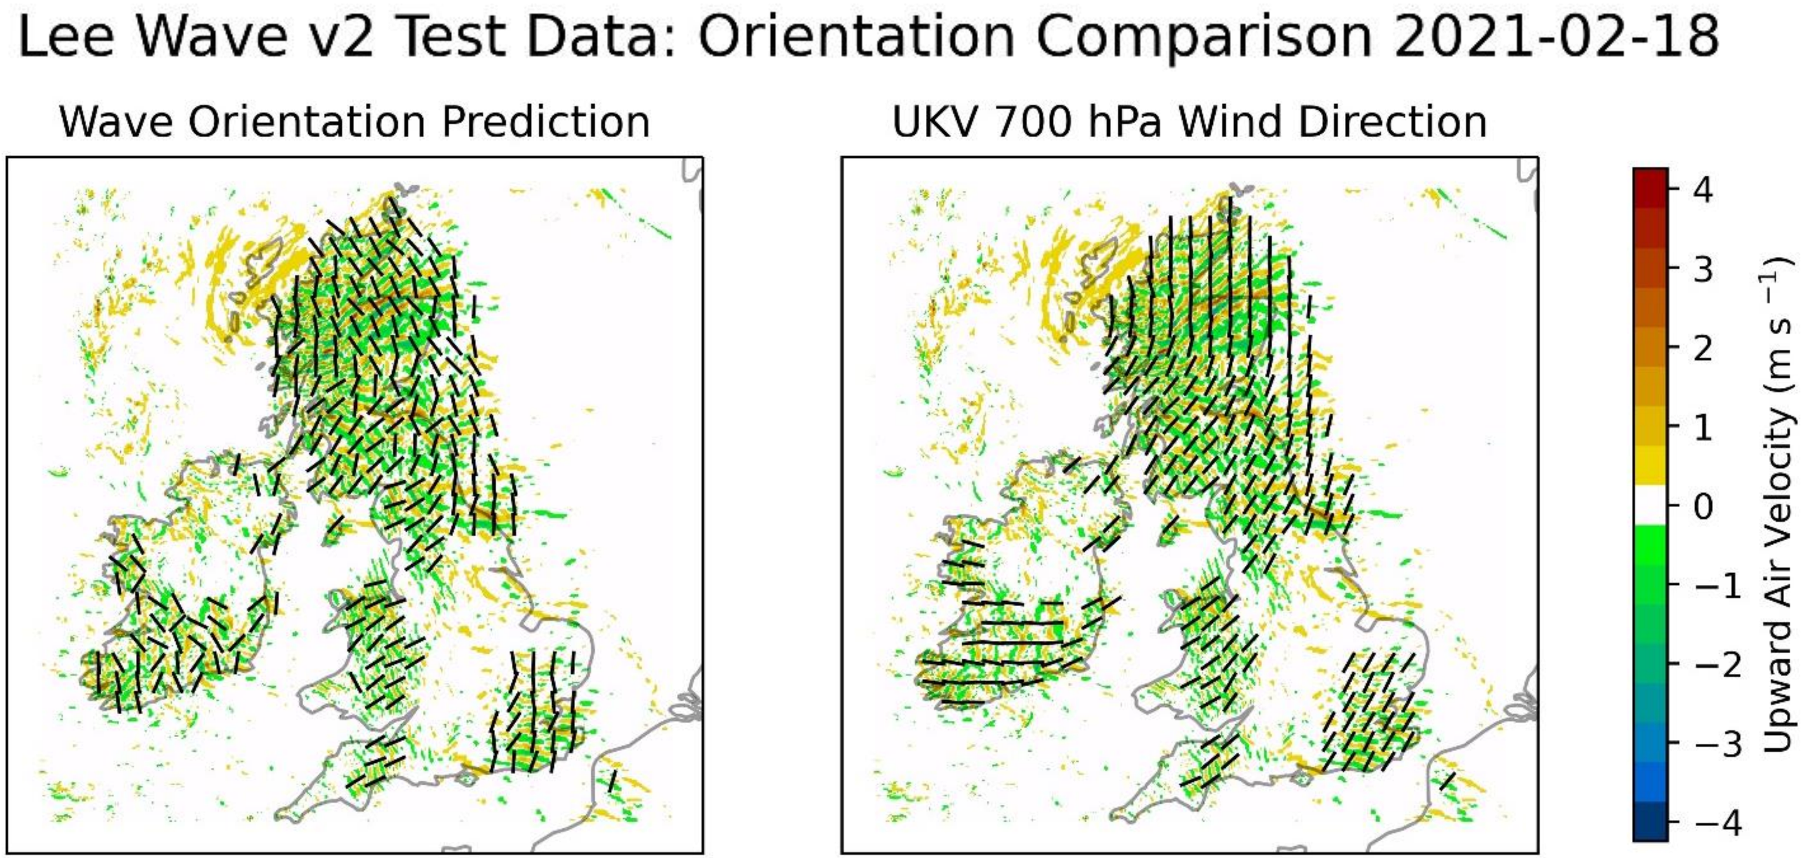

Summary

UNIVERSITY OF LEEDS

- Successfully trained a U-Net to segment lee waves into wave & no wave regions. It seems to have learned the difference between waves and convection.
- Looked into using Zooniverse to farm out labelling tasks to others, in an attempt to make better test data. With more volunteers giving more time, a larger training set could be created.
- Performed transfer learning to predict wave characteristics. Orientation looking promising when compared with other data, but wavelength needs some work to verify predictions.

Thanks to my supervisors, the Met Office orography group, the SciML group in ICAS and the Zooniverse volunteers. And thank you for listening to me ramble on.

References

- Durran, D. R., & Klemp, J. B. (1983). *Mountain Waves and Downslope Winds*. http://www.eumetrain.org/data/4/452/print_4.htm
- Durran, D. R. (2003). Lee Waves and Mountain Waves. In J. R. Holton, J. A. Curry, & J. A. Pyle (Eds.), *Encyclopedia of Atmospheric Sciences* (pp. 1161–1169). Elsevier Science. https://doi.org/10.1016/B0-12-227090-8/00202-5
- Sutherland, B. (2010). Internal waves in uniformly stratified fluid. In *Internal Gravity Waves* (pp. 141-212). Cambridge: Cambridge University Press. doi:10.1017/CBO9780511780318.004
- Holton, J. R., & Hakim, G. J. (2013). An Introduction to Dynamic Meteorology. In *An Introduction to Dynamic Meteorology* (5th ed., Vol. 9780123848). Elsevier. https://doi.org/10.1016/C2009-0-63394-8
- Sheridan, P., Vosper, S., & Brown, P. (2017). Mountain Waves in High Resolution Forecast Models: Automated Diagnostics of Wave Severity and Impact on Surface Winds. *Atmosphere*, *8*(12), 24. https://doi.org/10.3390/atmos8010024
- Teixeira, M. A. C. (2014). The physics of orographic gravity wave drag. *Frontiers in Physics*, *2*(July), 1–24. https://doi.org/10.3389/fphy.2014.00043

Lee Waves

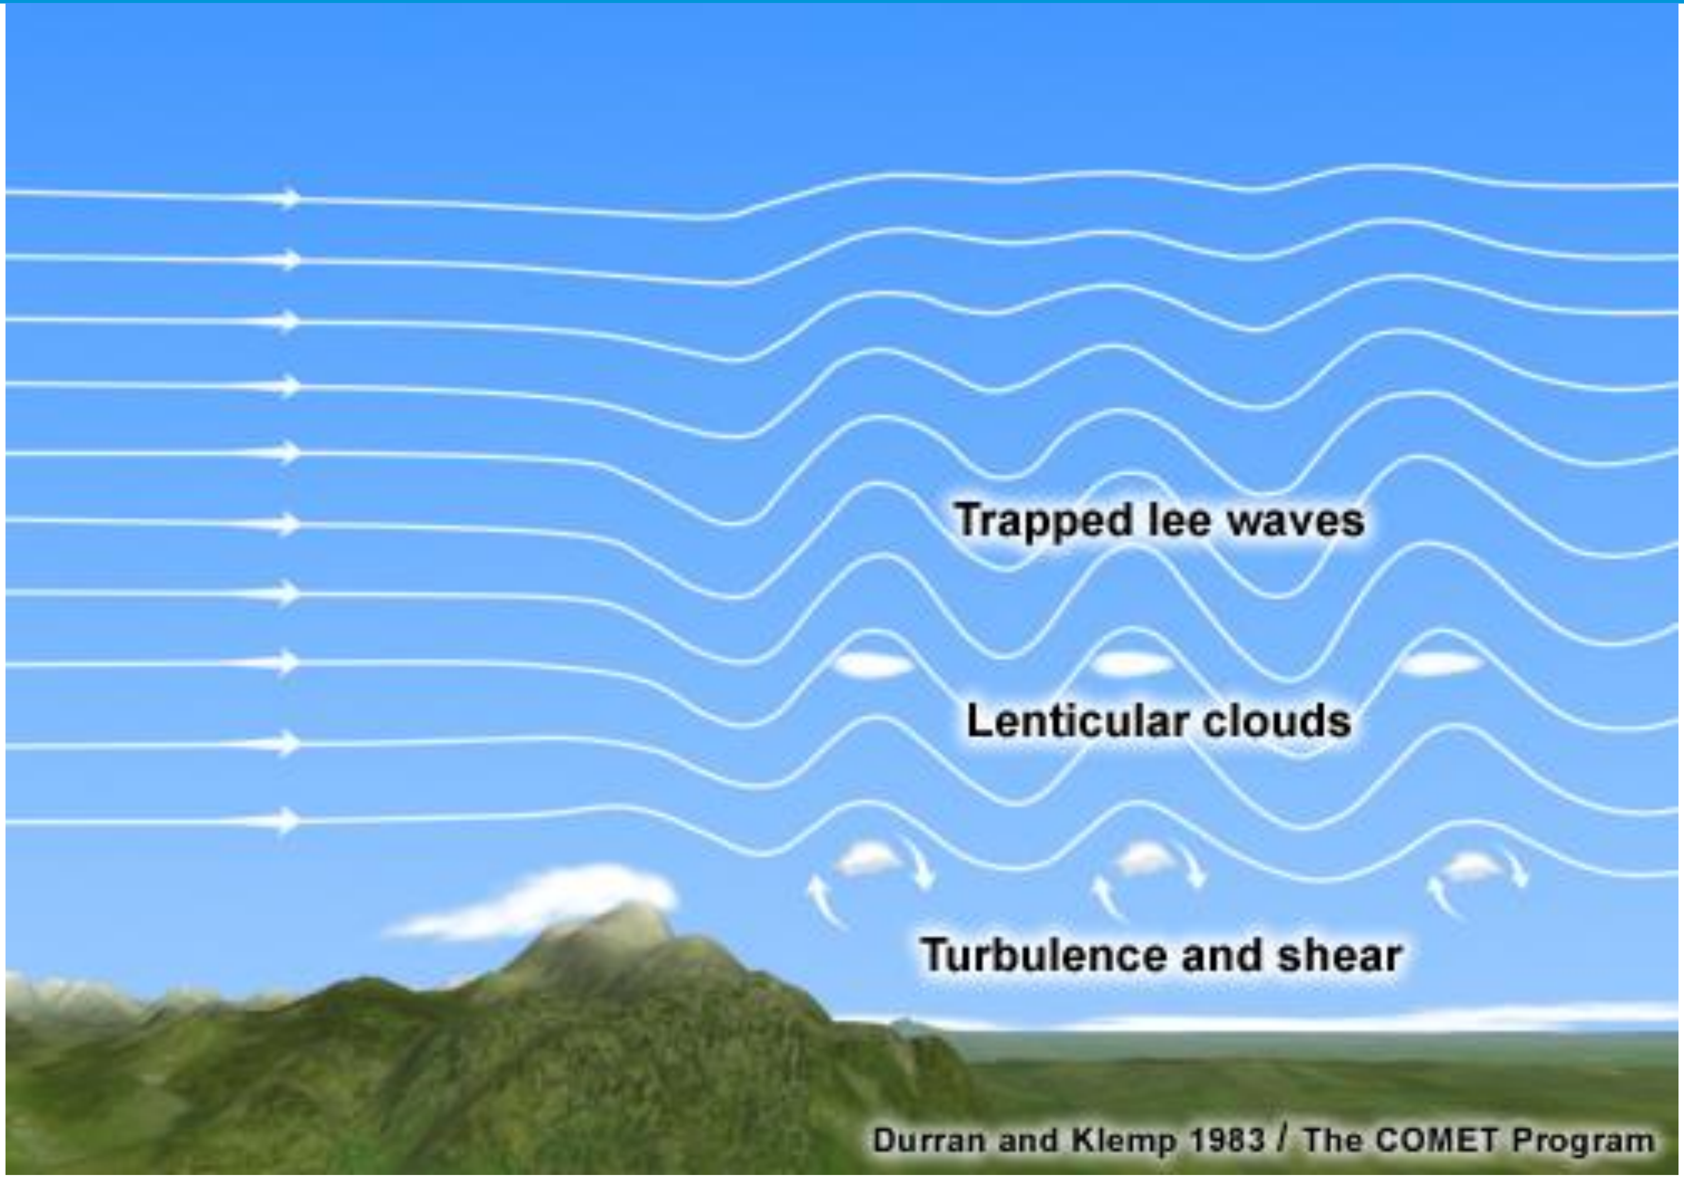
Trapped lee waves
Lenticular clouds
Turbulence and shear
Durran and Klemp 1983 / The COMET Program

Teixeira (2014)

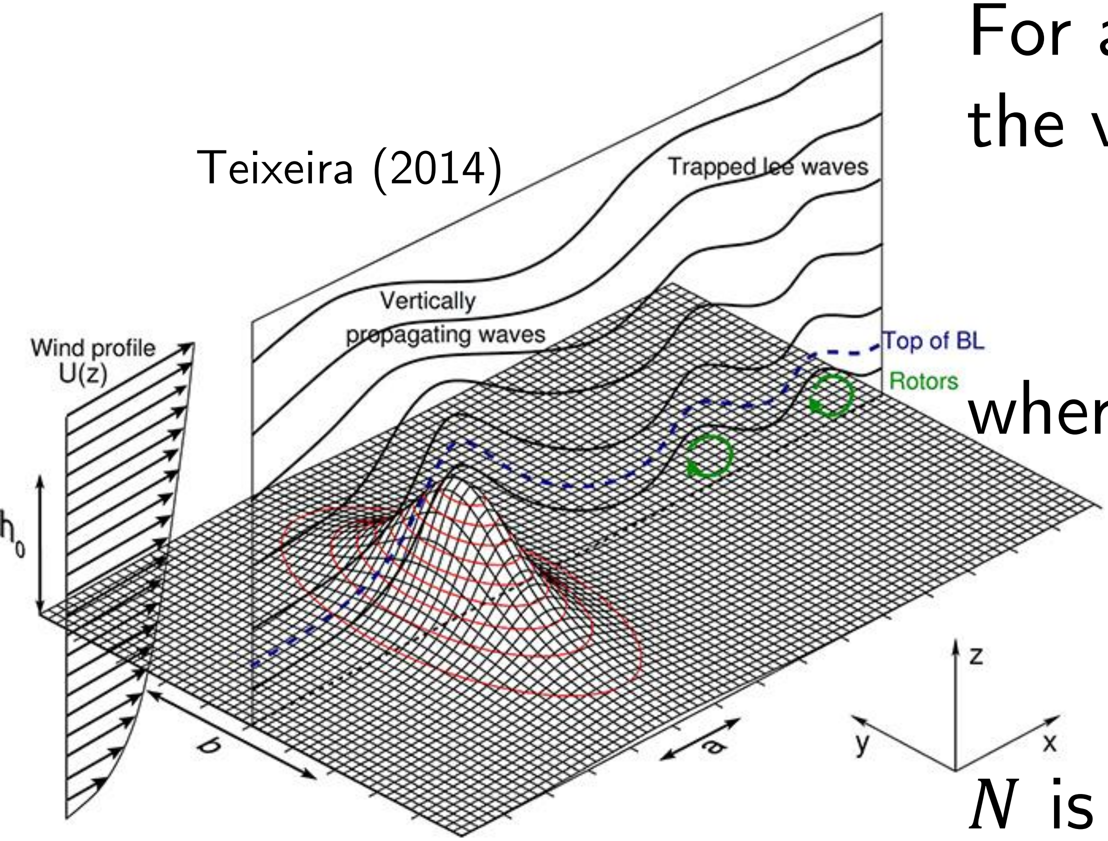

For a 2D flow over an obstacle, we can represent the vertical velocity $w(x,z)$ with:

$$\frac{d^2w}{dx^2} + \frac{d^2w}{dz^2} + l^2 w = 0 \qquad (**)$$

where l is the **Scorer Parameter:**

$$l^2(z) = \frac{N^2}{u^2} - \frac{1}{u}\frac{d^2u}{dz^2}$$

N is the Brunt–Väisälä frequency and u(z) is the horizontal wind profile as a function of height.

Durran (2003)

Simple case with a sinusoidal mountain

Simple example where u and N are constant with height: $\frac{d^2w}{dx^2} + \frac{d^2w}{dz^2} + l^2 w = 0\,(**)$

(**) has solutions $w(x,z) = \widetilde{w}_1(z)\cos kx + \widetilde{w}_2(z)\sin kx$ $\quad l^2(z) = \frac{N^2}{u^2} - \frac{1}{u}\frac{d^2u}{dz^2}$

Substituting into (**):

$$\frac{d^2\widetilde{w}_i}{dz^2} + (l^2 - k^2)\widetilde{w}_i = 0 \qquad i = 1,2$$

So we'll end up with different waves depending on the relationship between l^2 and k^2.

In this case if $l^2 - k^2 > 0$, waves decay horizontally and propagate vertically, while if $l^2 - k^2 < 0$, waves propagate horizontally.

Durran (2003)

Brunt–Väisälä frequency 1

UNIVERSITY OF LEEDS

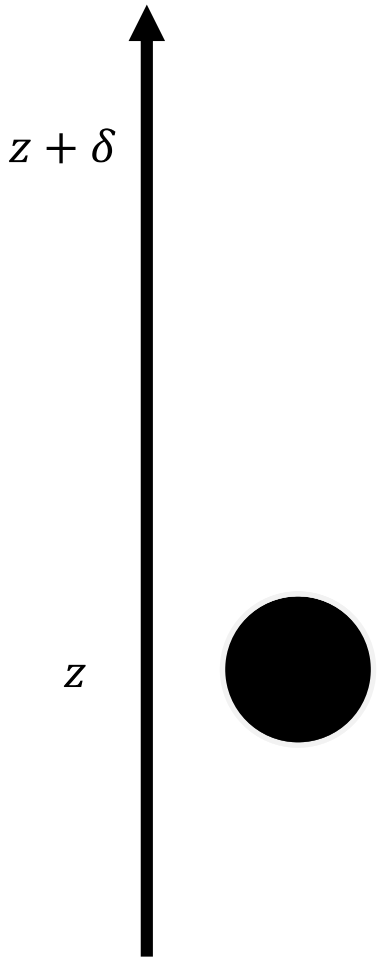

- Suppose we have a fluid parcel of density ρ_0 in a medium of density $\rho(z)$. It is nudged adiabatically from rest at some height z, to $z + \delta$.
- It is no longer at equilibrium, by Newton's 2nd law $(F = ma)$ the motion is:

Acceleration due to gravity

$$\rho_0 \frac{d^2\delta}{dt^2} = -g\underbrace{\big(\rho(z) - \rho(z+\delta)\big)}_{\text{Difference in density of the medium at height } z \text{ and } z+\delta}$$

Sutherland (2014)

Brunt–Väisälä frequency 2

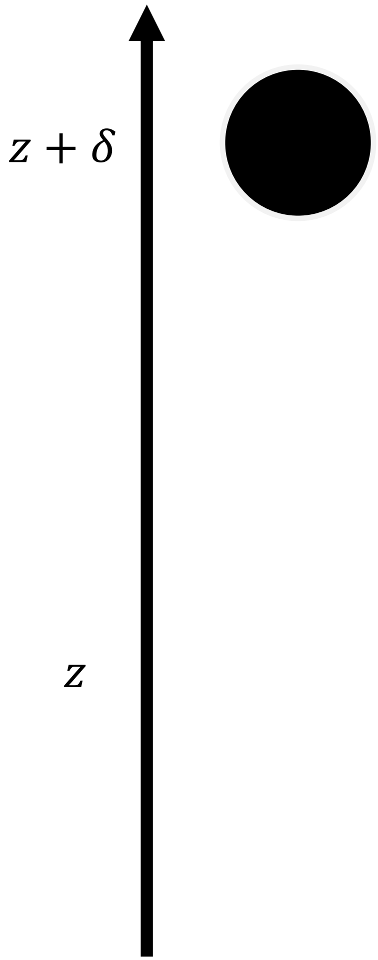

$$\rho_0 \frac{d^2\delta}{dt^2} = -g\big(\rho(z) - \rho(z+\delta)\big) \qquad (*)$$

Approximate $\rho(z+\delta) \approx \rho(z) + \frac{d\rho}{dz}\delta$, or $\rho(z) - \rho(z+\delta) \approx -\frac{d\rho}{dz}\delta$.

Substitute approximation into (*), $\rho_0 \frac{d^2\delta}{dt^2} = g\frac{d\rho}{dz}\delta$, or (tidying up):

$$\frac{d^2\delta}{dt^2} + N^2\delta = 0,$$

where

$$\boxed{N^2 = -\frac{g}{\rho_0}\frac{d\rho}{dz}}$$

N is the Brunt–Väisälä frequency (a measure of the frequency of parcel oscillation and used as a measure of stability).

Sutherland (2014)

Maths Summary

UNIVERSITY OF LEEDS

$$\text{Scorer Parameter: } l^2(z) = \frac{N^2}{u^2} - \frac{1}{u}\frac{d^2u}{dz^2}$$

- If u and N vary with height so that l decreases strongly with height:

  Favourable conditions for trapped lee waves

- l almost constant with height: vertically propagating waves
- l increases with height: unfavourable conditions for waves

Durran (2003)

3DVOM

UNIVERSITY OF LEEDS

- Existing Lee Wave Forecasting system 3DVOM:
  - Assumes uniform background flow, derived from one profile from the global model.
  - Linear model, simplified set of dynamical equations.
  - Limited area, only run over hilly areas

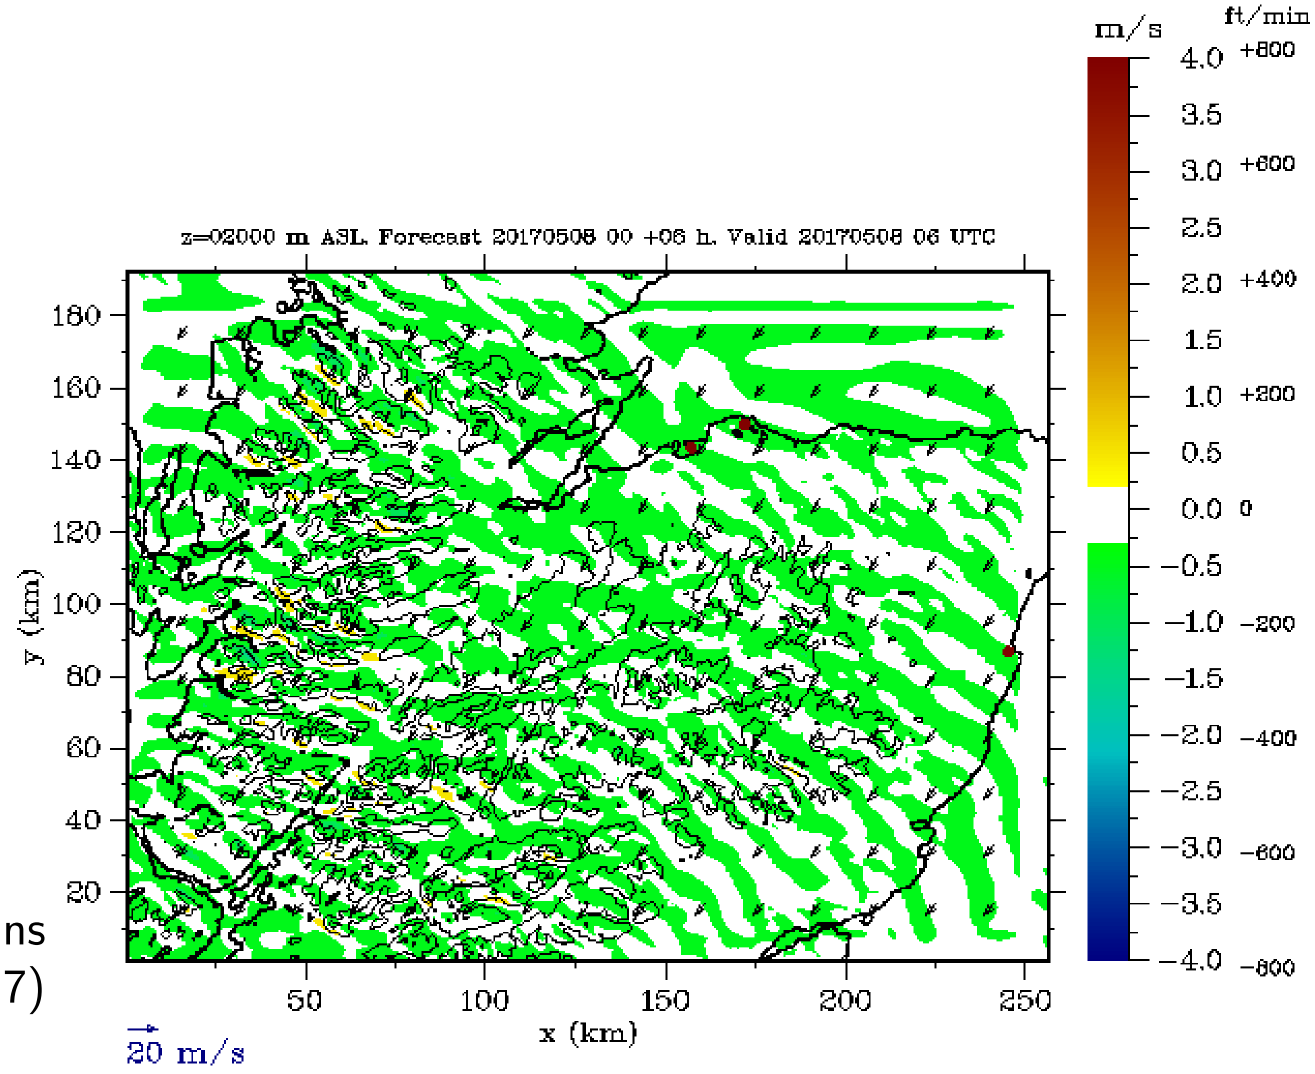

3DVOM Vertical Velocities at 2 km over the Grampians (Met Office, 2017)

Data Carpentry

- 363 examples of UKV vertical velocity output from January and February 2021 (335 training, 28 test), obtained from MASS using JASMIN.
- Selected 700 hPa as the height to use
- Cropped data to a 512 x 512 pixel grid (so 1024 x 1024 km)
- Custom tool used (2 slides ago) to create a mask for each data example of waves and not waves. Now we have a set of labelled data to train the U-Net to try to produce.

Training

UNIVERSITY OF LEEDS

- Split data up in to Training and Test sets. Test set not seen until the model is trained.
- Training set further split up into Train and Validation sets
- Weights and biases of model updated after each epoch based on loss function (in this case using built-in `FlattenedLoss of CrossEntropyLoss`).
- Continue training until validation loss starts increasing.
- Once the model is trained, we can then test it using our separate test set.

Results on the test set

UNIVERSITY OF LEEDS

- Verdict: not bad.

Model Name	Pixel Accuracy	IoU* score
U-Net v2	91%	0.68
U-Net v1	88%	0.61
Baseline (no waves anywhere)	83%	0.42

- Model is better than the baseline! Which is good.
- IoU score is a better metric than pixel accuracy:

We have far more events with no action than waves: test set 16% of the pixels are waves.

*Intersection over Union (also known as the Jaccard index).
If P = prediction & T = "truth", then $IoU = \frac{|P \cap T|}{|P \cup T|}$.

Some example results (model v2) on the test set

UNIVERSITY OF LEEDS

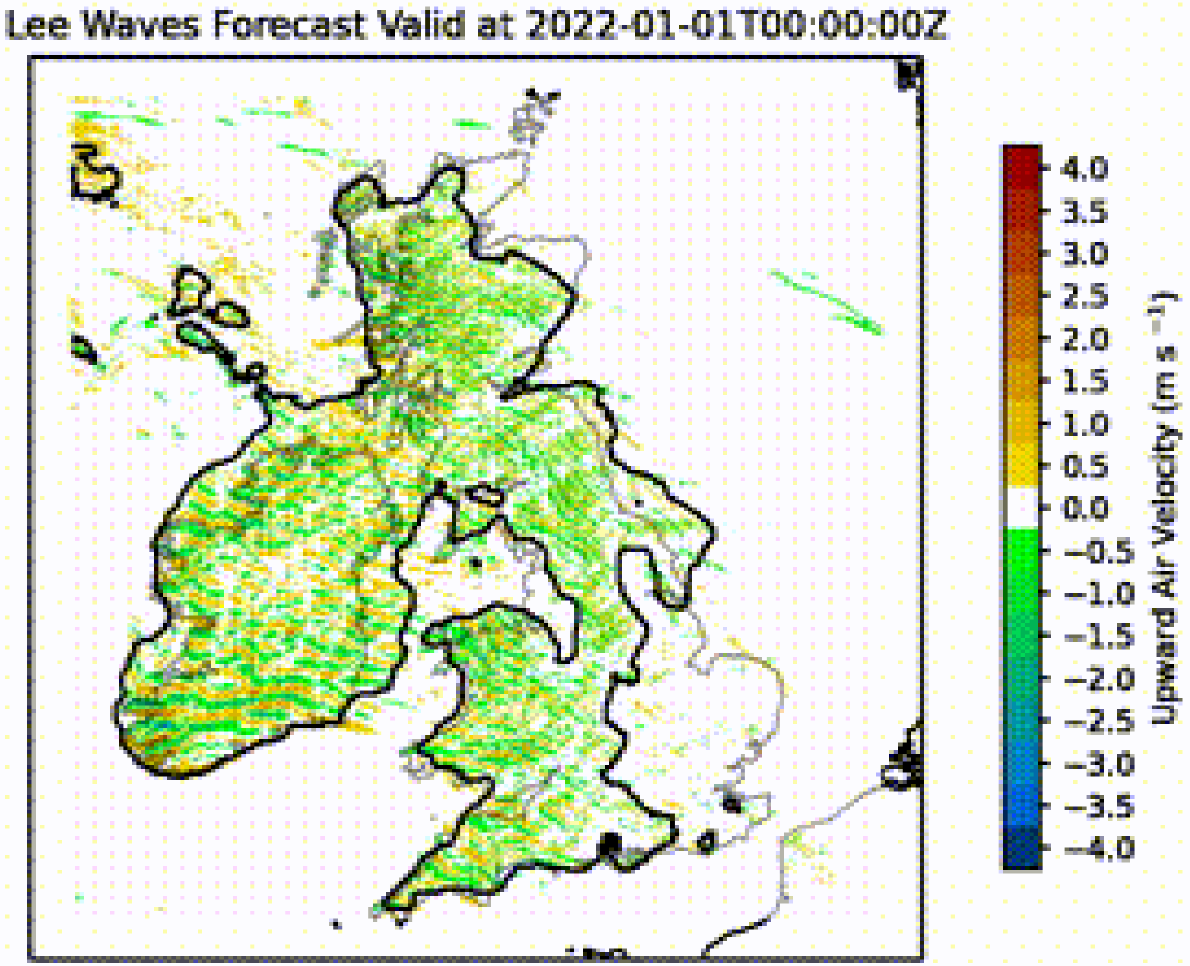

More examples: http://homepages.see.leeds.ac.uk/~mm16jdc/phd/lee_waves/

Amplitude too - using local maximum amplitude

UNIVERSITY OF LEEDS

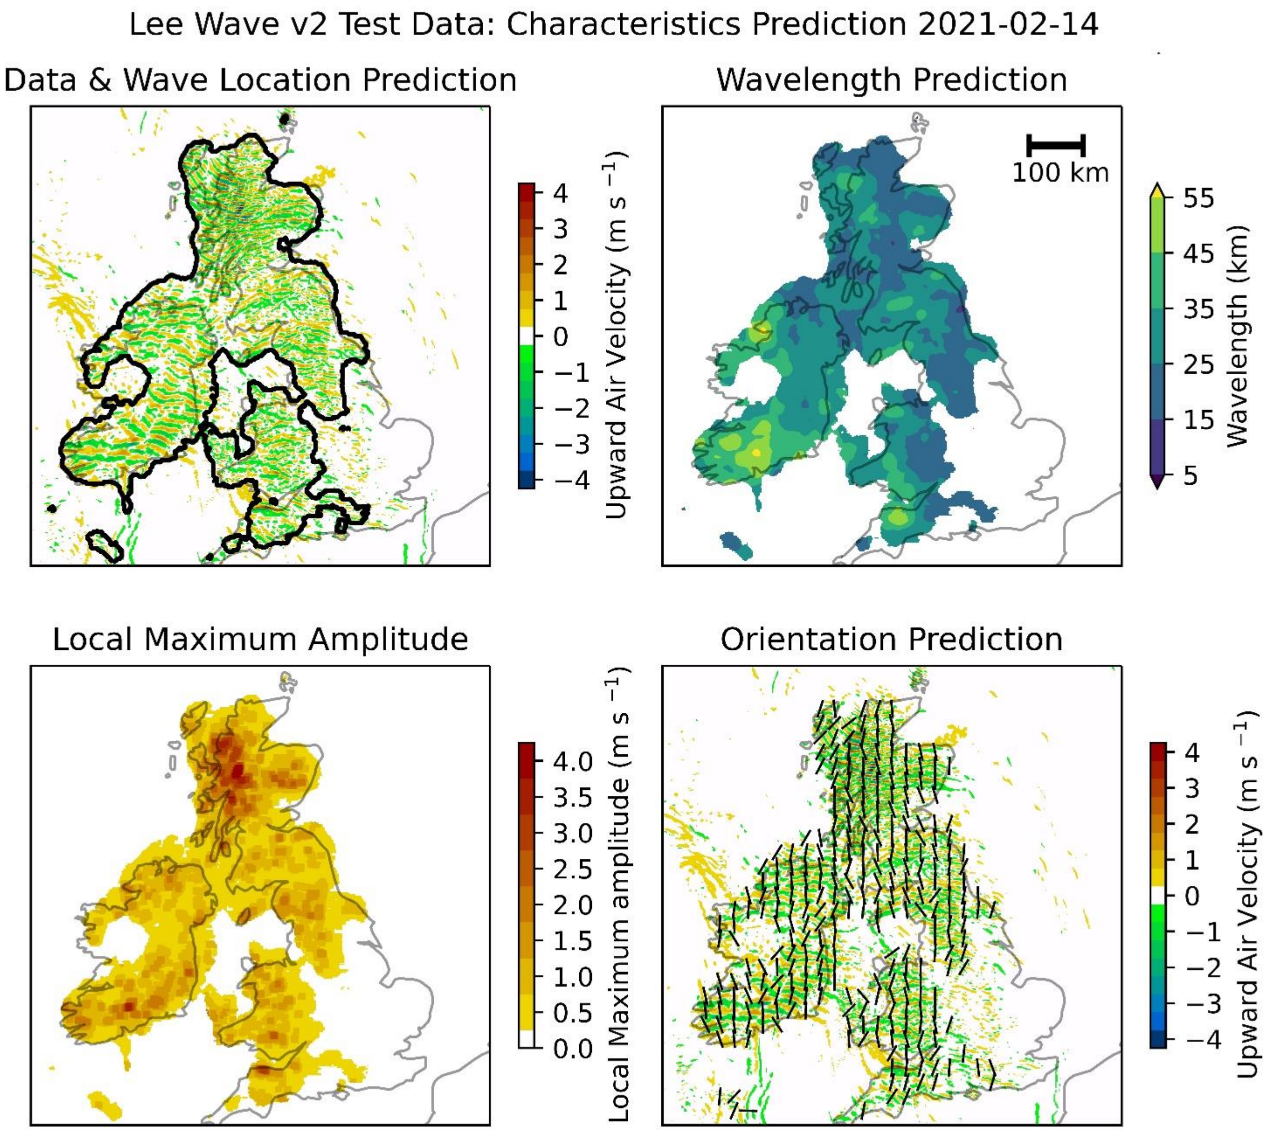

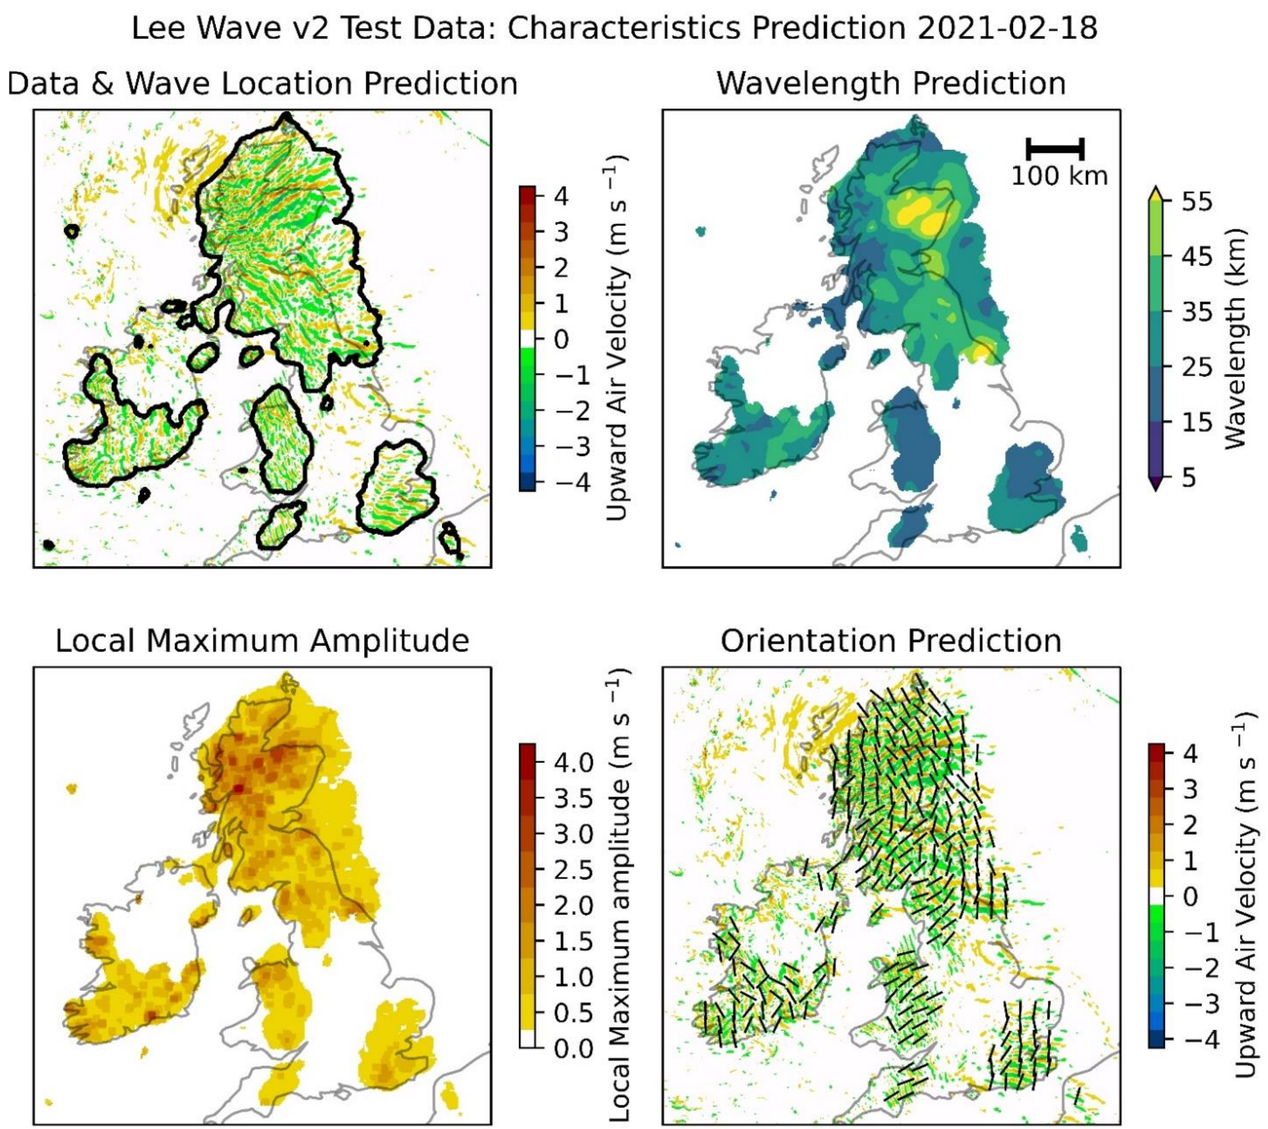

Mean Wave Orientation

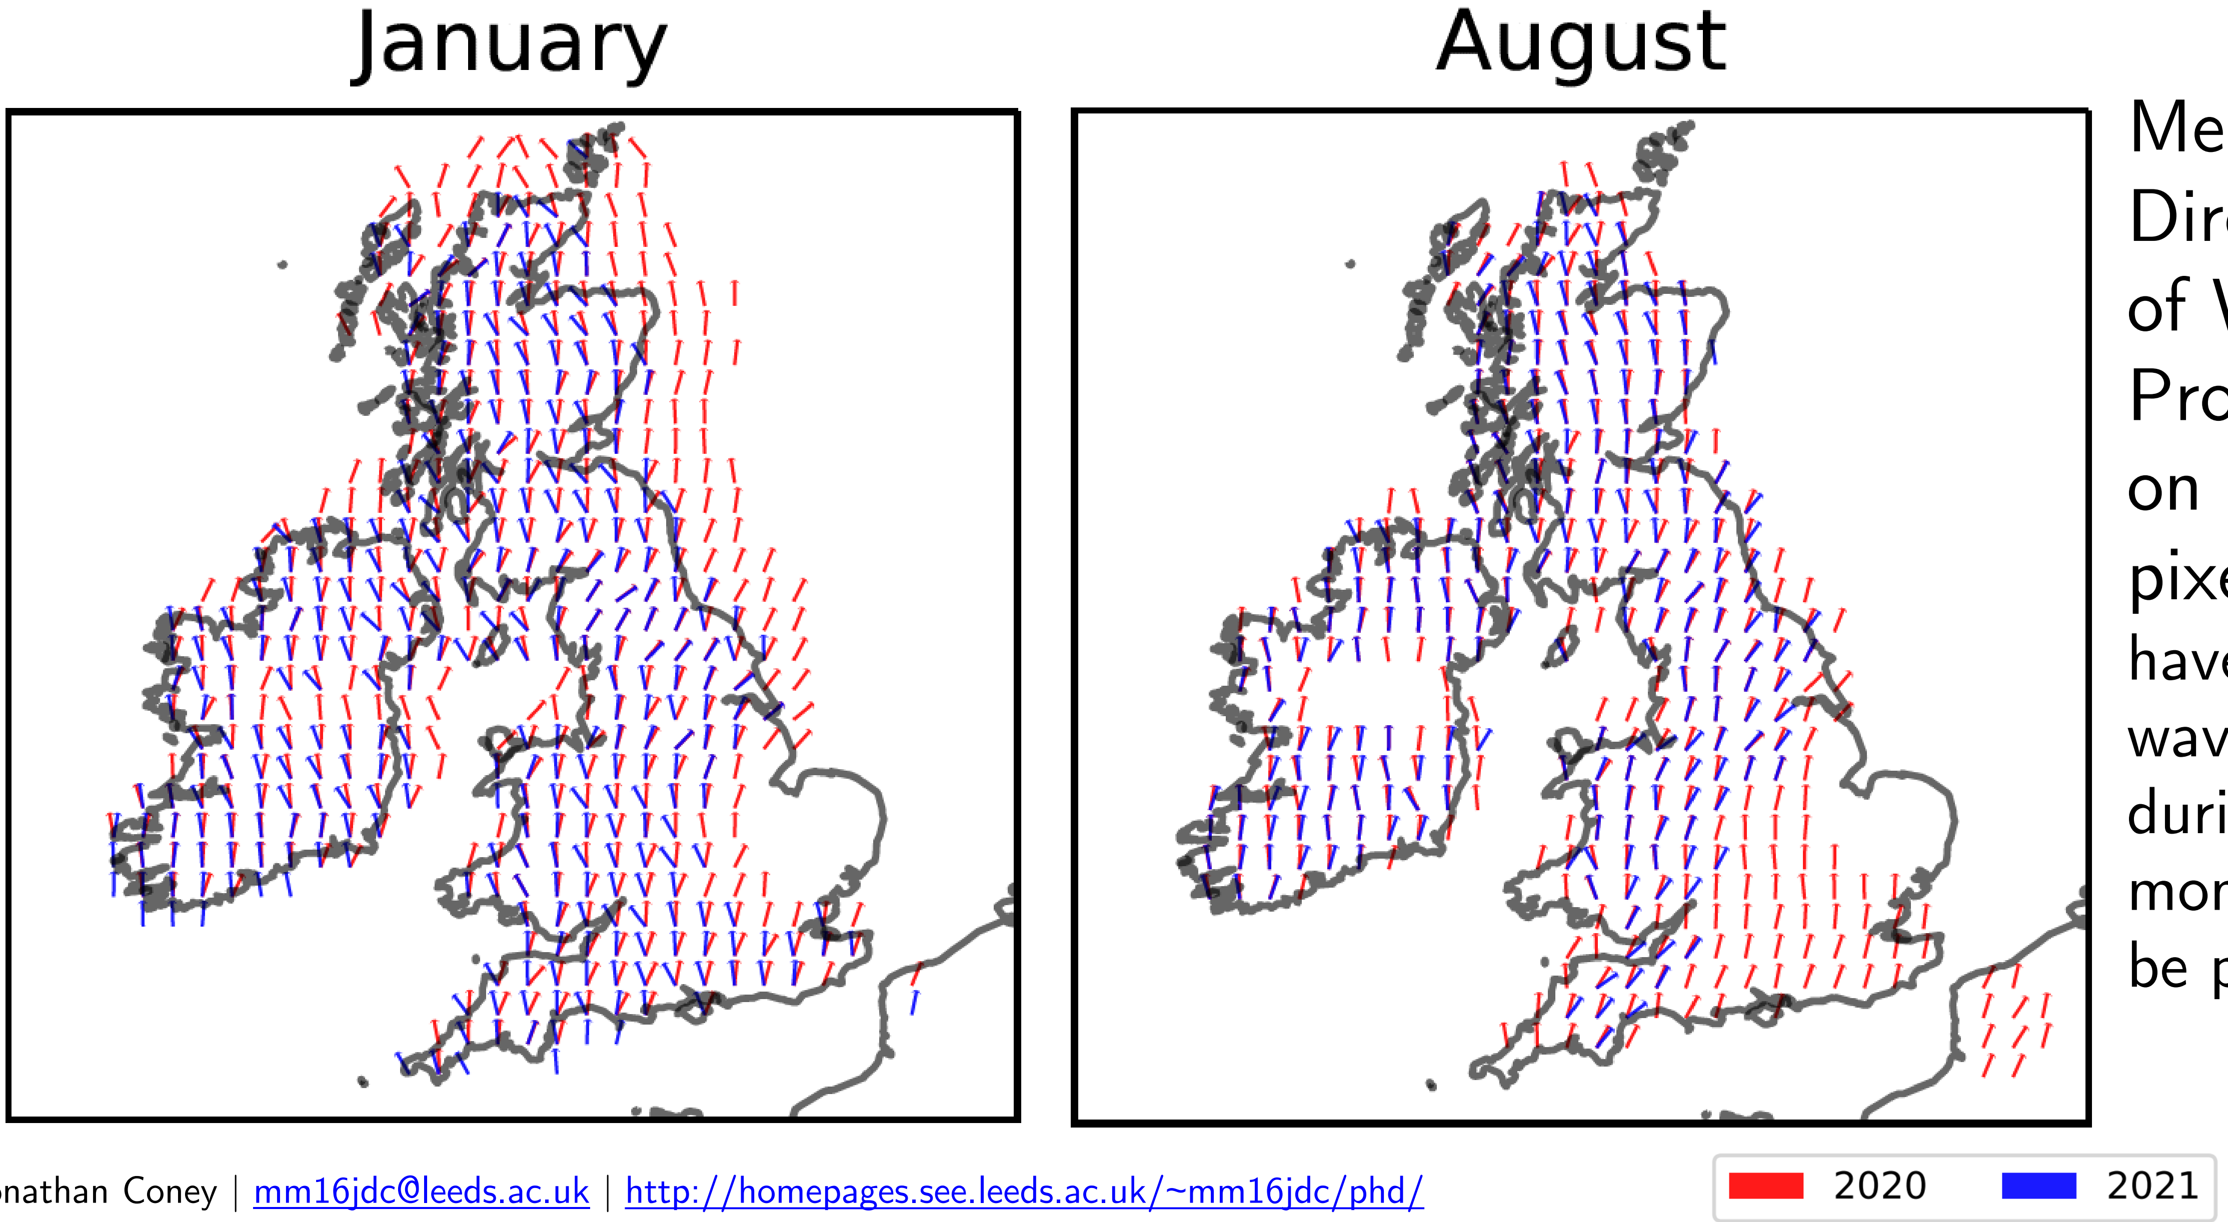

Mean Direction of Wave Propagation per pixel (must have >5% wave events during month to be plotted)

Wavelet technique on synthetic data: work needed

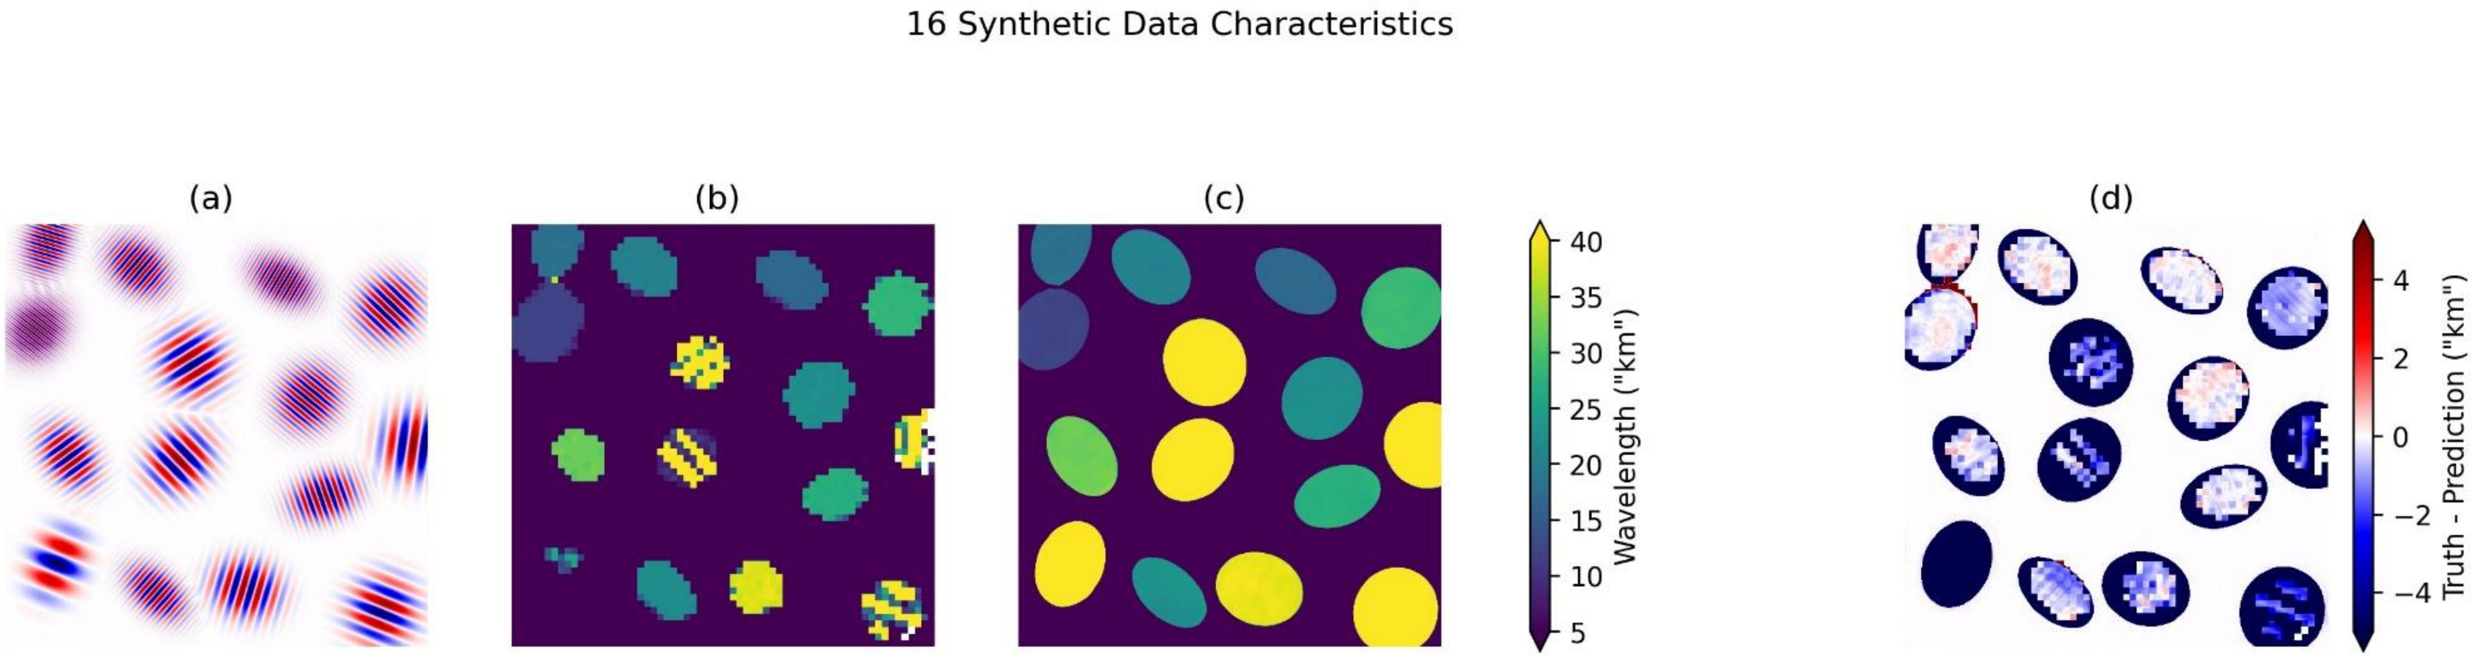

Wavelet attempt

- Peak detection using SciPy's `signal` library.
- First use ML learned orientation. Look either side of given pixel to get waves in that direction.
- Now get average wavelength from scikit peak function. Bit of a minefield as depends on sensitivity of peak function in scikit. Perhaps a link between segmentation model confidence and wavelength.

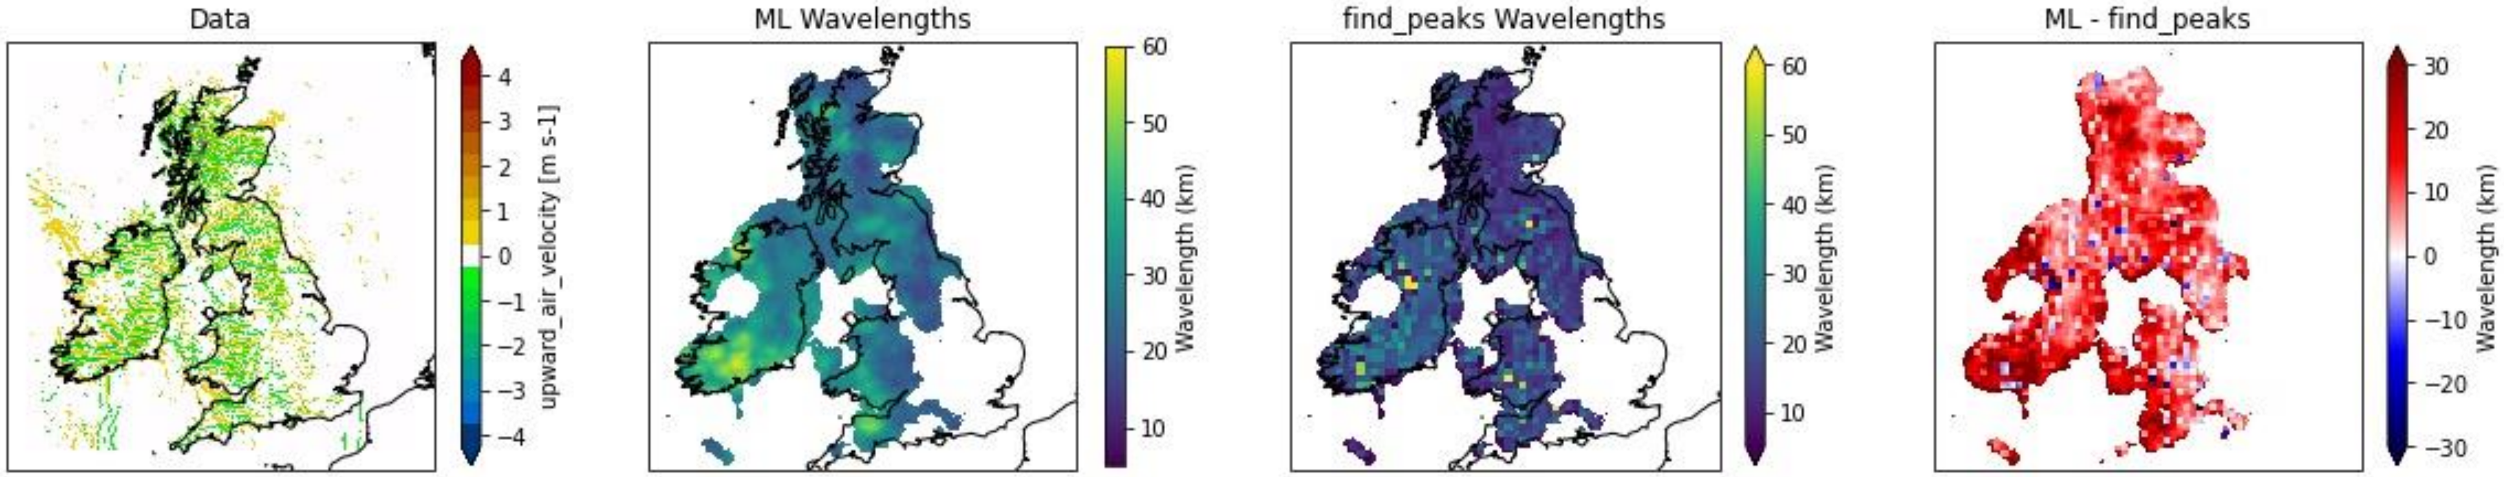

- We see more gravity waves in the data during the winter months.