A deep learning approach to extract internal tides scattered by geostrophic turbulence

Han Wang[1,1,1], Nicolas Grisouard[1,1,1], Hesam Salehipour[2,2,2], Alice Nuz[3,3,3], Michael Poon[1,1,1], and Aurelien L.S. Ponte[4,4,4]

[1]University of Toronto
[2]Autodesk Research
[3]New York University
[4]Ifremer

November 30, 2022

Abstract

A proper extraction of internal tidal signals is central to the interpretation of Sea Surface Height (SSH) data, yet challenging in upcoming satellite missions, where traditional harmonic analysis may break down at finer observed spatial scales known to contain significant wave-mean interactions. However, the wide swaths featured in such satellite missions render SSH snapshots that are spatially two-dimensional, which allows us to treat the tidal extraction as an image translation problem. We design and train a conditional Generative Adversarial Network, which, given a snapshot of raw SSH from an idealized numerical eddying simulation, generates a snapshot of the embedded tidal component. We test it on synthetic data whose dynamical regimes are different from the data provided during training. Despite the diversity and complexity of data, it accurately extracts tidal components in most individual snapshots considered and reproduces physically meaningful statistical properties.

Posted on 30 Nov 2022 — CC-BY-NC-ND 4 — https://doi.org/10.1002/essoar.10508849/v3 — This a preprint and has not been peer reviewed. Data may be preliminary.

A deep learning approach to extract internal tides scattered by geostrophic turbulence

Han Wang[1*], Nicolas Grisouard[1], Hesam Salehipour[2], Alice Nuz[1†], Michael Poon[1‡], and Aurélien L. Ponte[3]

[1]Department of Physics, University of Toronto, Ontario, Canada
[2]Autodesk Research, Ontario, Canada
[3]Ifremer, Plouzané, France

Key Points:

- A deep conditional Generative Adversarial Network is trained to extract tidal components in SSH snapshots from an idealized model.
- The network can extract tidal signals accurately in a snapshot whose underlying dynamics are different from training data.
- Performance of the network degrades when extracting tidal signals entangled with higher turbulence energies.

*Current address, School of Mathematics, University of Edinburgh, Edinburgh, UK
†Current address, Tandon School of Engineering, New York University, New York, USA
‡Current address, Department of Astronomy and Astrophysics, University of Toronto, Ontario, Canada

Corresponding author: Han Wang, hannnwangus@gmail.com

Abstract

Extraction of internal tidal (IT) signals is central to the interpretation of Sea Surface Height (SSH) data. The increased spatial resolution of future wide-swath satellite missions poses a challenge for traditional harmonic analysis, due to prominent and unsteady wave-mean interactions at finer scales. However, the wide swaths will also produce spatially two-dimensional SSH snapshots, which allows us to treat IT extraction as an image translation problem for the first time. We design and train TITE, a conditional Generative Adversarial Network, which, given a snapshot of raw SSH from an idealized eddying simulation, generates a snapshot of the embedded IT component. We test it on data whose dynamical regimes are different from the data provided during training. Despite the diversity and complexity of data, it accurately extracts ITs in most individual snapshots considered and reproduces physically meaningful statistical properties. Predictably, TITE's performance decreases with the intensity of the turbulent flow.

Plain Language Summary

Wide-swath satellite observations of Sea Surface Height (SSH) data at high spatial resolutions will be available in abundance thanks to advances of instrumental technologies. Embedded in the observed SSH are internal tides, a physical process that plays a crucial role in ocean circulation. As they are entangled with background currents and eddies, such tidal signals are challenging to extract. On the other hand, the wide satellite swaths provide new opportunities as they allow us to regard the observations as spatially two-dimensional. Here we treat the tidal extraction solely as an image translation problem. We train TITE, a deep neural net that, given a snapshot of a raw SSH signal, produces a "fake" snapshot of the tidal SSH signal that is meant to reproduce the original. The data we use in this article is generated by idealized numerical simulations. Once adapted to realistic data, the network has the potential to become a new tidal extraction tool for satellite observations. More broadly, successes in our experiments can inspire other applications of generative networks to disentangle dynamical components in data where classical analysis may fail.

1 Introduction

Internal tides (hereafter "ITs") are inertia-gravity waves generated by the ocean large-scale (barotropic) tidal currents flowing over submarine topographic features. They play large roles in deep/upper ocean mixing (Munk & Wunsch, 1998; Whalen et al., 2020; Lien & Gregg, 2001; Lahaye et al., 2019; Garrett, 2003), altimetric inference of balanced flows (Fu & Ferrari, 2008), and detection of ocean temperature changes (Zhao, 2016). For decades, IT extraction has been conducted via harmonic analysis (Munk & Hasselmann, 1964; Munk & Cartwright, 1966; Cartwright & Ray, 1990) . Current altimetry has a typical spatial resolution of $O(100)$ km (Ballarotta et al., 2019), which is sufficient to retrieve mode-1 and some of the mode-2 IT wavelengths of semidiurnal tides, along with the dominant turbulent balanced motions (hereafter "TBMs"; see Ray & Zaron, 2011). At these scales, the coupling between ITs and TBMs is usually weak and therefore substantial portions of the ITs are phase-locked with astronomical forcings (Egbert & Ray, 2000; Zaron, 2017). Harmonic fits can retrieve such phase-locked signals, even if the time sampling is longer than the tidal period, provided that the time series is long enough to diagnose the aliased lower frequencies (Carrere et al., 2021).

When ITs become incoherent (i.e., their phase shift with respect to their generating astronomical tide becomes time-dependent) due to scatterings from TBMs or modulations of mean density profiles, techniques such as complex demodulation (Munk et al., 1965; Colosi & Munk, 2006), least-square fits (Cartwright & Ray, 1990), Lagrangian filtering (Shakespeare & Hogg, 2017) or combination with in-situ data (Geoffroy & Nycander, 2022) are useful if well-sampled time series are available. For satellite observa-

tions, this is sometimes indeed the case, especially with multi-satellite observations (Zaron, 2015).

However, such techniques yield increasingly poor performance as temporal sampling rates decrease. According to several studies (Ray & Zaron, 2011; Zaron, 2015, 2017; Nelson et al., 2019), between 15% and 45% of mode-1 semidurnal tidal variances can not be identified via conventional harmonic analysis, with more recent estimates reporting larger fractions of incoherence. Starkly, using 17 years' and 20 years' global satellite observations respectively, Shriver et al. (2012) and Zhao et al. (2016) find apparent disappearances of mode-1 semidurnal tidal signals in equatorial Pacific regions, which Buijsman et al. (2017) attribute to equatorial incoherence rather than dissipation, as echoed in the dominance of incoherent tidal variance reported in Zaron (2017) and Nelson et al. (2019). Improvements of incoherent tidal extraction are hence required to better identify spatial distributions of tidal dissipation, which is crucial for predicting tidal mixing (Whalen et al., 2020). Such challenges will be exacerbated in the next generation of satellite altimetry, in particular the Surface Water Ocean Topography (SWOT) mission. In addition to SWOT's coarse sampling period of up to 21 days (Morrow et al., 2019), another complication arises from the improved spatial resolution of a few tens of km in wavelength (Morrow et al., 2019). At these smaller spatial scales, higher-mode ITs tend to become more incoherent (Ray & Zaron, 2011; Dunphy & Lamb, 2014; Shriver et al., 2012; Carrere et al., 2021) due to stronger couplings with the TBMs linked to the increased relative vorticity magnitudes (Bühler, 2014; Zaron & Egbert, 2014).

Attempts to circumvent time series information and focus on the spatial domain date back to at least Ray & Zaron (2011). New opportunities will arise in future altimeters from the availability of data along wide swaths (two 50 km-wide swaths, 20 km apart for SWOT) as opposed to current linear tracks. As a result, they will produce spatially two-dimensional (2D) images and it will be possible to regard the extraction of IT signals as an operation on 2D snapshots. Recently proposed methods rely on exploiting distinct wavenumber spectral signatures of TBMs and internal waves (Torres et al., 2019), empirical modal mapping methods (Egbert & Erofeeva, 2021), dynamical relations to surface density fields (Ponte et al., 2017), or on data assimilation techniques (Metref et al., 2020; Le Guillou et al., 2021). Our goal is to propose an alternative method based on a deep-learning, image processing method.

Indeed, several theoretical arguments indicate that valuable information is contained in the spatial distribution of SSH field. Statistically, internal tides are concentrated around tidal wavenumbers determined by the Sturm-Liouville problem (Gill & Adrian, 1982; Ray & Zaron, 2011; Zaron, 2017). Also, in the TBM component, kinetic energy and enstrophy transfer rates derived from SSH satisfy properties from classical turbulence theories (Khatri et al., 2018). Alternatively, ray tracing or triad resonance theories provide deterministic constraints in highly idealized cases (Savva & Vanneste, 2018; Ward & Dewar, 2010). To probe whether spatial information alone is enough to extract ITs in a deterministic fashion, we propose to regard the IT extraction solely as an image-to-image translation problem and to develop what we call the "Toronto Internal Tide Emulator" (TITE), a deep convolutional neural network that extracts the SSH signature induced by IT from a raw, instantaneous SSH map via pattern recognition. In general, we find TITE to perform well in most SSH snapshots generated from a set of idealized simulations.

2 Methods

2.1 Idealized data supporting TITE's development

Data to support TITE's development are snapshots from a set of idealized numerical simulations, where mode-1 ITs are forced at a fixed tidal period T (12 hours) to prop-

agate through TBMs created by a baroclinically unstable jet (Ponte & Klein, 2015; Ponte et al., 2020). The SSH signatures of TBMs in these simulations are generally larger than those induced by ITs, and exhibit a significant overlap in spatial scales at $O(100)$ km with ITs. Spatial filtering is thus difficult, an issue that is also faced by satellite altimetry in oceanic regions such as the Gulf Stream or Drake Passage, where powerful TBMs exist (Rocha et al., 2016; Richman et al., 2012).

We run the model under five different initial meridional density contrasts. With increasing contrast, the baroclinic jet becomes more unstable and creates a more vigorous baroclinic eddy field. The spectra induced by these eddies follow a geostrophic turbulence law (Ponte & Klein, 2015; Charney, 1971), and are thus identified as TBMs. In ascending order of stationary surface kinetic energy levels of TBM (hereafter referred to as "turbulence levels"), we label the five simulations as T1 to T5. The mean normalized vorticities (absolute values of surface vertical vorticities normalized by the local Coriolis frequency) over the jet width increase from 0.06 to 0.14 from T1 to T5. We refer to Ponte & Klein (2015) for details on the numerical setup, and recapitulate the relevant information in Text S1 in the Supporting Information (SI).

IT snapshots are computed online via harmonic fits over time series that are $2T$ long and sampled every 300 seconds, or $T/144$. For simplicity, we only study $\eta_{\cos}^{(\mathrm{sim})}$, the cosine component of ITs from the simulations, defined in the same way as Dunphy et al. (2017):

$$\eta_{\cos}^{(\mathrm{sim})}(x,y,t) = \frac{1}{T}\int_{t-2T}^{t} \eta(x,y,t')\cos\left(\frac{2\pi}{T}t'\right)\mathrm{d}t', \tag{1}$$

where x, y are the zonal and meridional coordinates, respectively, and η denotes raw SSH. For each snapshot, we cut out three square panels covering three fixed latitudinal bands, labeled as "down-jet", "mid-jet" and "up-jet", as illustrated in Fig. 1. Here, the term "up-jet" denotes locations closer to the tidal forcing regions. One hundred snapshots are captured every $4T$ for each simulation in T1-5, resulting in 1500 pairs of $\left\{\eta,\ \eta_{\cos}^{(\mathrm{sim})}\right\}$ panels (5 runs, 3 latitudinal bands, and 100 snapshots) altogether.

2.2 Deep-learning algorithm designed to extract tidal signals

During the design of the TITE runs, we implicitly make four assumptions: (1) there is abundant spatial information, (2) all snapshots are statistically independent from each other in time, (3) a raw SSH functionally determines its IT component, but properties of the functional dependence are unknown, and (4) there exists abundant data where ITs are already extracted from the raw SSH. We discuss these assumptions at the end of this article.

TITE is based on a conditional Generative Adversarial Network (hereafter referred to as "cGAN"). As the name implies, a cGAN consists of two parts, namely, a conditional generator (hereafter "generator") that learns how to manufacture a "fake" image that's conditioned on an "input image", and a discriminator that tries to determine if an image is "genuine" (i.e., paired to the input image in the training data), or fake (i.e., created by the generator). Either part is on its own a convolutional neural network, and during training, the two parts compete against each other to co-evolve (Mirza & Osindero, 2014; Goodfellow et al., 2014). We denote the cosine IT panels generated from TITE as $\eta_{\cos}^{(\mathrm{gen})}$; following our notations, the *input* image would be η, the *genuine* image would be $\eta_{\cos}^{(\mathrm{sim})}$, and the *fake* image would be $\eta_{\cos}^{(\mathrm{gen})}$. As reflected in this general workflow, during training, other than the paired panels, no further information is given to TITE.

The particular cGAN TITE is derived from is called "pix2pix" (Isola et al., 2017), applications of which range from artistic creations (ml4a, 2017) to scientific problems such as remote sensing image classifications (Lebedev et al., 2018). Our codes are adapted from a TensorFlow Tutorial (Tensorflow, n.d.). We refer to Text S6 and S4 in the SI for

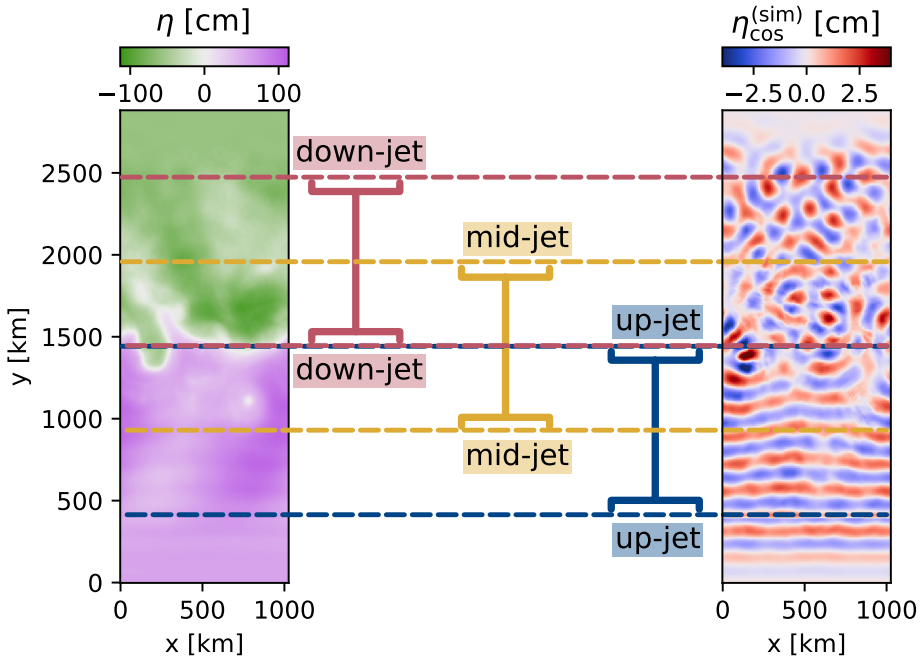

Figure 1. The "down-jet", "mid-jet" and "up-jet" bands illustrated over a snapshot of η (left) and $\eta_{\cos}^{(\mathrm{sim})}$ (right), sampled from T3 at day 2120. The "mid-jet" band is centred around the baroclinic jet. ITs are forced to the south of "up-jet" bands, and as the ITs propagates northward and loses coherence due to interactions with the TBM, the $\eta_{\cos}^{(\mathrm{sim})}$ patterns are less reminiscent of plane waves in the "down-jet" band than in the "up-jet" band.

more details on our design considerations, and provide the code link in the Acknowledgements. Here, we summarize the essential features.

The generator and the discriminator have around 10^4 and 2000 convolutional layers respectively. The considerable number of model parameters makes TITE a black box, as in the case of many deep learning algorithms.

Prior to each epoch, training images are randomly reshuffled in time, cropped, flipped, and rotated. Here, an epoch means the number of computations it takes for the cGAN to iterate over all data in the training set once. The random cropping, rotation and flipping are intended to introduce challenges that roughly mimic realistic situations where we don't have a priori knowledge of the observer's orientation/location about IT generation sites, direction of propagation, and exact boundary conditions. By randomly reshuffling in time, we enforce that every panel pair at every snapshot in the simulation be sequentially independent from the others. This means that any temporal information in the simulations is unknown to the pix2pix kernel, in line with our assumption (2) made previously in this section.

As the fully convolutional U-Net structure inherited from pix2pix in the generator can be applied to images of arbitrary sizes in principle, when producing Movies S1-5 in the SI, we directly apply the trained TITE onto rectangular input images, even though TITE is trained on square images illustrated in Fig. 1. This versatility on the shapes of input images would be useful for along-swath satellite products.

We systematically run our code with TensorFlow 2.3.0 under Python 3.7. One hundred training epochs with 960 pairs of $\left\{\eta, \eta_{\cos}^{(\mathrm{sim})}\right\}$ in the training set take about 1.5 hours with a NVIDIA GP100 GPU. For all the TITE runs in the article, we choose to present the results after 600 training epochs. .

2.3 Division of data to training, testing and validation sets

As a first check on TITE's performance, we randomly select 20% of all 1500 pairs of $\left\{\eta, \eta_{\cos}^{(\mathrm{sim})}\right\}$ panels from T1-5 to form a so-called validation set, and use the rest as the training set. During training, TITE has access to all pairs of $\left\{\eta, \eta_{\cos}^{(\mathrm{sim})}\right\}$ in the training set, but none from the validation set. After 600 epochs, the training phase is over, and we apply the trained TITE into snapshots in the validation set. The mean correlation between $\eta_{\cos}^{(\mathrm{sim})}$ and $\eta_{\cos}^{(\mathrm{gen})}$ in the validation set turns out to be 0.85, which suggests that the generated $\eta_{\cos}^{(\mathrm{gen})}$ reasonably resemble the ground truths $\eta_{\cos}^{(\mathrm{sim})}$. However, under this division, the training set contains turbulence levels that are statistically similar to the validation set on which the trained TITE is applied, and the good correlation factors could be caused by *overfitting*. Here,"overfitting" refers to a commonly accepted definition in machine-learning contexts (Dietterich, 1995): a model is said to overfit when it tries to fit the training data so closely that it does not generalize well to new data. To address this possibility, we challenge TITE to extract $\eta_{\cos}^{(\mathrm{sim})}$ signals linked to a *different* turbulence level as those employed for its training.

Specifically, in what we refer to as the "ET1 run", we reserve a *test* set, which contains all 300 pairs of panels from the simulation T1 and *none* from T2, T3, T4 or T5. Among the remaining panels from T2-5, we randomly select 80% pairs for the training set, and reserve the other 20% for the validation set. The validation and test sets are both inaccessible to TITE during training, but crucially, in terms of average turbulence levels, the training set is similar to the validation set, yet *different* from the test set. This procedure is designed to limit, or at least detect, the occurrence of overfitting. Similarly, we carry out ET2-5 runs, following the same logic, where the test sets are panels from the simulations T2-5 respectively.

3 Performance of TITE

As a graphical illustration, in Movie S1 in the SI, we re-order all the shuffled test instances of ET1 in time. The reconstructed temporal information appears remarkable, considering that the snapshots were randomly shuffled at the time of their generation and hence that the temporal evolution of these images was unknown to TITE. However, as turbulence level in test sets increases from ET2 to ET5, the evolution of $\eta_{\cos}^{(\mathrm{gen})}$ bears less and less resemblance with $\eta_{\cos}^{(\mathrm{sim})}$ (Movies S2-5 in the SI). To find the underlying causes, in this section, we evaluate the performance of TITE with several statistical metrics and discuss the causes of relatively decreased performance when they arise. All metrics are computed using standard methods and detailed in Text S5 in the SI .

We first investigate how close $\eta_{\cos}^{(\mathrm{gen})}$ is to the ground truth $\eta_{\cos}^{(\mathrm{sim})}$ by measuring the correlation between the two, as in Torres et al. (2018). The mean correlation factors averaged over all test instances in the ET1-5 runs are 0.91, 0.89, 0.83, 0.80 and 0.70, respectively. The highly correlated predictions of TITE in the test set in ET1-4 are especially interesting, as turbulence levels of the test set are different from that of the training set. There is however a relatively sharper drop (from 0.80 to 0.70) in the mean correlation from ET4 to ET5 in the test sets. Such a sharp drop can also be observed in the histograms shown in the SI Fig. S7.

A few relevant snapshots from test sets are illustrated in Fig. 2. The first row presents the snapshot with the highest correlation among all test sets in ET1-5. It belongs to ET1

and has a correlation factor of 0.95. The second and third row present the snapshot with the lowest correlation among all test sets in ET5 and ET4, respectively. Judging by the correlations, the worst case in ET4 has a correlation factor of 0.68, which is a significant improvement over the correlation factor of 0.49 in the worst case in ET5. This agrees with visual comparisons between the second and third rows in Fig. 2. This observation, together with the sharp drop of mean correlation factors in ET5 noted above, suggest a possible categoric difference between ET5 and ET1-4.

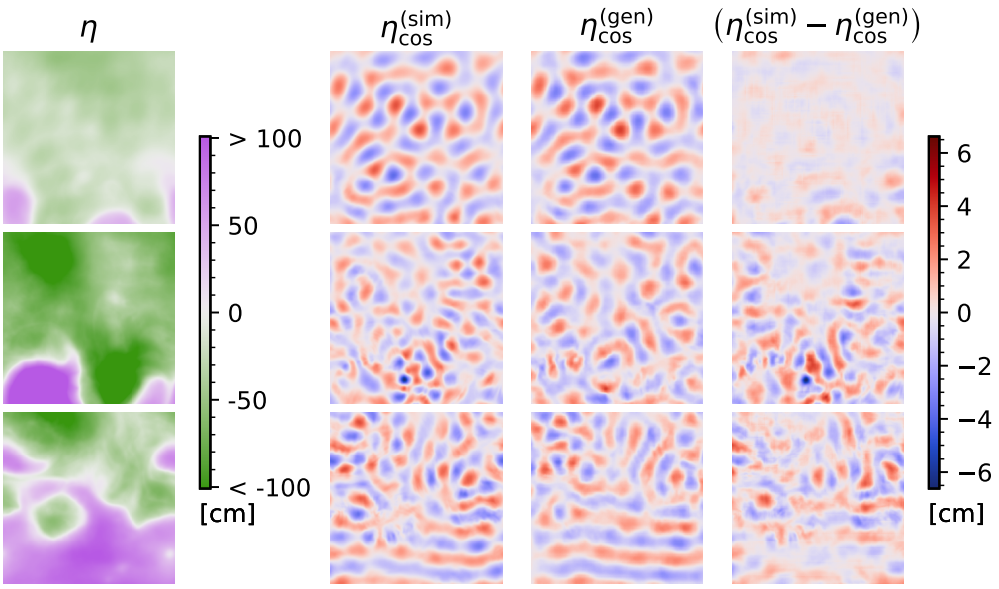

Figure 2. A few examples of individual tests. Top row: snapshot with the highest correlation among all sets (belongs to ET1). Middle and bottom rows: same for the lowest correlations among ET1-4, and ET1-5, respectively (belong to ET4 and 5, respectively). For legibility, we omit spatial axis labels, noting that the first, second, and third rows correspond to down-jet, down-jet, and mid-jet bands respectively; see Fig. 1 for their locations.

To gain more insight about the relative failures in ET5, we conduct a spectral analysis that focuses on comparing ET4 and ET5. The wavenumber spectra for the down-jet and up-jet bands are computed separately for $\eta_{\cos}^{(\mathrm{sim})}$ and $\eta_{\cos}^{(\mathrm{gen})}$ in the test set of ET4 and ET5, and presented in Fig. 3. The spectra for mid-jet bands are omitted for readability here and attached in Text S2 in the SI .

Prominent bumps appear near the wavenumbers corresponding to mode-1 tidal wavelengths in all the spectra of $\eta_{\cos}^{(\mathrm{sim})}$ (Solid lines in Fig. 3). These bumps are somewhat broad, and their locations are noticeably different between the down-jet and up-jet bands. This is expected, as the density profiles and the Coriolis parameter both vary with latitude, which modulates the mode-1 tidal wavelength. Such variations can also be found in satellite observations (Ray & Zaron, 2011).

A qualitative difference can be found in the spectral behaviours between ET4 and ET5. In ET4, the locations of spectral bumps in the $\eta_{\cos}^{(\mathrm{gen})}$ spectra vary between the down-jet and up-jet bands, in a manner such that they *closely* overlap with bumps of the $\eta_{\cos}^{(\mathrm{sim})}$

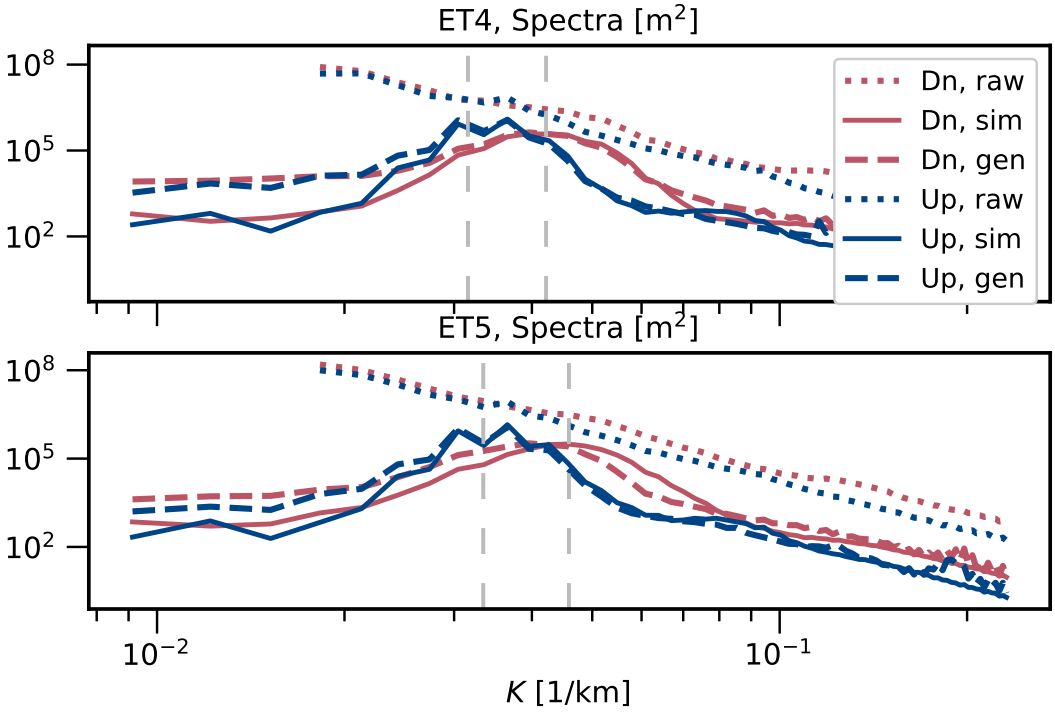

Figure 3. Spectra for the down-jet and up-jet bands in ET4 and ET5 test set. In the legends "Dn", "Up" denote the down-jet and up-jet bands respectively. "raw", "sim", and "gen" denote spectra computed from panels of $\eta, \eta_{\cos}^{(\mathrm{sim})}$, and $\eta_{\cos}^{(\mathrm{gen})}$, respectively. "$K$" denotes the horizontal wavenumber magnitude. The vertical dashed lines mark the largest and smallest mode-1 tidal wavenumbers over the simulation domain at initial time. Raw spectra higher than $2 \times 10^8 \mathrm{m}^2$ at large scales are omitted. Higher wavenumbers are omitted.

spectra at both bands. This implies that in the ET4 run, the trained TITE identifies the dominant wavelength even as it varies. In other words, TITE can identify patterns at varying spatial scales. However, in the ET5 run, the $\eta_{\cos}^{(\mathrm{gen})}$ spectra *fail* to trace the location of the bumps in the down-jet bands (dashed and solid red lines in lower panel of Fig. 3). The performance in up-jet bands appears as good as ET4, which may be attributed to the fact that the mode-1 tidal wavelengths to the south of the jets are the same in all five simulations.

What causes the decrease of performance in ET5? We can formulate two possible hypotheses: (i) the inherent difficulty of predicting increasingly complex patterns at higher turbulence levels, and (ii) overfitting (see section 2.3 for its definition). Overall, we believe factor (i) affects TITE's performance *more* than factor (ii), as we now argue.

In support of hypothesis (i), stronger scatterings of ITs from TBMs induce more longitudinal variations as well as small-scale features in the IT components due to increased incoherence. In addition, the tidal wavelengths vary more latitudinally due to increased density gradients, which increases the diversity of dominant spatial scales of IT signals across the domain and time. Both factors add complexities to the η and $\eta_{\cos}^{(\mathrm{sim})}$ patternsAs the turbulence levels increase from T1 to T5, the increased pattern complexities in the test set pose more challenges to our algorithm, which explains the decreased performance in ET5 in correlations and spectral behaviors.

The difficulty associated with vigorous turbulence levels is also reflected in the relatively worse performance of TITE in the mid-jet bands centered around the turbulence. We find that within each of ET1-5, the up-jet bands have a higher mean correlation than the mid-jet bands, and as the turbulence level increases, this difference gets more pronounced (See the histograms in Fig. S7 of the SI for quantitative behaviors). Spectral behaviors degrade too, as detailed in Text S2 in the SI.

Turning our attention over to hypothesis (ii), in our experiments, overfitting may provide a partial explanation for the performance degradation from ET4 to ET5. As expected from section 2.1, the kinetic energy and normalized vorticity for the TBM and IT all increase from T1 to T5 (with the former increasing from 0.04 m^2/s^2 to 0.20 m^2/s^2). In terms of these dynamical metrics, the training set of ET5 is less diverse compared to, say, the training set of ET4 that spans a wider range of these metrics. Therefore, in ET5, TITE may be fitting closely for turbulence levels in T1-4, and fails at T5 due to its mismatch with T1-4, whereas in ET4, both T3 and T5 are included during training, and a combination of T3 and T5 may impart some knowledge about T4 during testing.

However, several behaviours of TITE are inconsistent with the overfitting hypothesis. First, if overfitting was the only factor, then during the test phase, TITE should perform worse in ET1 than in ET2, as the range of dynamical metrics in ET1 is narrower than in ET2. This does not turn out to be the case. In fact, the ET1 run produces the best mean correlation in the test set among ET1-5. Second, compared to test sets, validation sets are closer to training sets, and an overfitting model should perform better in the validation sets than in the test sets. However, the mean correlations produced by TITE in the test sets are higher than the validation sets by about 0.05 in ET1 and ET2.

On the other hand, hypothesis (i) provides better explanations for these disagreements. In the ET1 test set, the turbulence levels are the lowest, and TITE performs well despite the possible impacts from overfitting. In ET1 and ET2, the test data are at a lower turbulence level than the validation data, and TITE generates better predictions. Based on hypothesis (i), to improve TITE's performance at higher turbulence levels, a possible solution is to feed TITE with "overly turbulent" simulations during training, as motivated by the better performances we have found in ET4, as compared to ET5.

4 Conclusions and Discussions

We designed a novel technique based on a deep neural network algorithm to extract internal tides that are entangled with geostrophic turbulence. We trained and validated TITE using randomly shuffled simulation snapshots that were categorically different from the dynamic regime of the testing data. Due to lack of temporal information, strong overlapping in spatial scales of ITs and TBMs, and incoherence, harmonic fits or space filtering alone could not extract tidal signals accurately in the testing data sets, and yet in most test cases, TITE can still (1) extract IT signals that agree well with ground truths in a deterministic sense, and (2) capture the dominant tidal energy in the wavenumber spectra, even when it varies temporally and latitudinally. When TITE does not perform as well, the main cause seems to be the high complexities of the patterns linked to stronger turbulent motions. Overall, we believe that this work provides a fresh angle on how to disentangle dynamical components from two-dimensional data via a deep learning approach. Some discussions are offered below.

Although we make no claim about TITE or cGANs in general as being the best possible algorithms to specifically achieve our goal, we found it superior to other deep learning methods we investigated, which include several types of decision trees regressors, long short-term memory networks, and U-Net structures without a discriminator. We did not attempt to optimize model parameters such as numbers of layers or learning rates, among others. Improvements on pix2pix such as that of Park et al. (2019) could

also improve our current implementation. Moreover, as mentioned in section 3, the generated images contain spurious signals outside the dominant tidal bump, which remains to be resolved. We leave these as thoughts for future work.

So far, we have only applied TITE to the idealized simulations T1-5 with a single baroclinic jet and single tidal frequency, simplistic boundary conditions, flat topography, an absence of atmospheric forcings. As an ongoing work, we are investigating the effects of including snapshots from a global ocean GCM.

We used correlation factors and wavenumber spectra as main metrics in our evaluations of TITE's performance, similar to several other disentanglement efforts (Torres et al., 2019; Ponte et al., 2017). However, no direct physical meanings are attached to correlation factors, while no quantitative metric assessing spectral behaviours are widely adopted for practical applications. Our community has yet to widely adopt a skill metric to assess the accuracy of 2D tidal extraction.

With SWOT in mind, we may reassess the four assumptions stated in the first paragraph of section 2.2. All images used in this work have a 4 km horizontal resolution that resolves the tides adequately, addressing assumption (1). In preparation for satellite data that will suffer from measurement error and more limited resolutions, we may coarse-grain and augment the training data with the type of error expected in SWOT (Gaultier et al., 2016) to investigate their impacts. We could then rigorously compare TITE with other IT models (Carrere et al., 2021). Assumption (2), motivated by the incoherence of ITs and the relatively long sampling intervals of SWOT, is satisfied by the design of the TITE architecture, and by the frequent random shuffling of snapshots during training. However, complete statistical independence between ITs and TBMs can be overly strict for several reasons, ranging from a higher temporal sampling at high latitude, to the possibility of "filling in the time gaps" with other sources of data such as those from assimilated models or in-situ instruments (d'Ovidio et al., 2019). From the overall satisfactory performance of TITE, assumption (3) appears to be satisfied in our simulation outputs, due to simulation design choices such as a perfectly harmonic incoming IT signal, or simple boundary conditions. Under more realistic configurations, a functional dependence might not be guaranteed. On the other hand, assumption (3) can also be too strong a constraint, considering the known theoretical knowledge of IT/TBM interactions (Savva & Vanneste, 2018; Ward & Dewar, 2010). Assumption (4) relies on the premise that there will be pre-processed training data (presumably from highly skilled model outputs) that mimic the dynamics to be sampled by SWOT. Productions of such data are receiving significant attention within the modelling communities (Zaron & Rocha, 2018; Rocha et al., 2016; Arbic et al., 2010; Savage et al., 2017). Overall, to make TITE eventually applicable to SWOT and other satellite missions in the future, more work is required, especially in coordination with different communities and perhaps combination with complementary methods.

Acknowledgments

We thank two anonymous reviewers for numerous constructive advice. H.W. and N.G. acknowledge the financial support of the Canadian Space Agency [14SUSWOTTO] and of the Natural Sciences and Engineering Research Council of Canada (NSERC) [RGPIN-2015-03684]. M.P. acknowledges support from the Summer Undergraduate Research Program of the Department of Physics at the University of Toronto. A.L.P. was supported by ANR project number 17-CE01-0006-01 and by project DIEGO of the CNES/TOSCA SWOT program. This work was made possible by the facilities of the Shared Hierarchical Academic Research Computing Network (SHARCNET: www.sharcnet.ca) and Compute/Calcul Canada. We have benefited from discussions with Laure Zanna, Brian K. Arbic, Michael W. Lever, Jason E. Summers, Jacques Vanneste and Ilia A. Iakovlev.

Simulation results corresponding to snapshots in T1-5 used in this study to train and test TITE are published on Scholars Portal Dataverse (Ponte et al., 2020). Codes defining the architecture of TITE are available at: `https://github.com/hannnwang/Pix2PixTITEexamples`

References

Arbic, B. K., Wallcraft, A. J., & Metzger, E. J. (2010). Concurrent simulation of the eddying general circulation and tides in a global ocean model. *Ocean Modelling*, *32*(3-4), 175–187.

Ballarotta, M., Ubelmann, C., Pujol, M.-I., Taburet, G., Fournier, F., Legeais, J.-F., ... others (2019). On the resolutions of ocean altimetry maps. *Ocean Science*, *15*(4), 1091–1109.

Bühler, O. (2014). *Waves and Mean Flows* (2nd ed.). Cambridge: Cambridge University Press. Retrieved from `http://ebooks.cambridge.org/ref/id/CBO9781107478701` doi: 10.1017/CBO9781107478701

Buijsman, M. C., Arbic, B. K., Richman, J. G., Shriver, J. F., Wallcraft, A. J., & Zamudio, L. (2017). Semidiurnal internal tide incoherence in the equatorial pacific. *Journal of Geophysical Research: Oceans*, *122*(7), 5286–5305.

Carrere, L., Arbic, B. K., Dushaw, B., Egbert, G., Erofeeva, S., Lyard, F., ... Picot, N. (2021, 1). Accuracy assessment of global internal-tide models using satellite altimetry. *Ocean Science*, *17*, 147-180. doi: 10.5194/os-17-147-2021

Cartwright, D. E., & Ray, R. (1990). Oceanic tides from geosat altimetry. *Journal of Geophysical Research: Oceans*, *95*(C3), 3069–3090.

Charney, J. G. (1971). Geostrophic turbulence. *Journal of the Atmospheric Sciences*, *28*(6), 1087–1095.

Colosi, J. A., & Munk, W. (2006). Tales of the venerable honolulu tide gauge. *Journal of physical oceanography*, *36*(6), 967–996.

Dietterich, T. (1995). Overfitting and undercomputing in machine learning. *ACM computing surveys (CSUR)*, *27*(3), 326–327.

Dunphy, M., & Lamb, K. G. (2014). Focusing and vertical mode scattering of the first mode internal tide by mesoscale eddy interaction. *Journal of Geophysical Research: Oceans*, *119*(1), 523–536.

Dunphy, M., Ponte, A. L., Klein, P., & Le Gentil, S. (2017). Low-mode internal tide propagation in a turbulent eddy field. *Journal of Physical Oceanography*, *47*(3), 649–665.

d'Ovidio, F., Pascual, A., Wang, J., Doglioli, A. M., Jing, Z., Moreau, S., ... others (2019). Frontiers in fine-scale in situ studies: Opportunities during the swot fast sampling phase. *Frontiers in Marine Science*, *6*, 168.

Egbert, G., & Erofeeva, S. (2021). An approach to empirical mapping of incoherent internal tides with altimetry data. *Geophysical Research Letters*, e2021GL095863.

Egbert, G., & Ray, R. (2000). Significant dissipation of tidal energy in the deep ocean inferred from satellite altimeter data. *Nature*, *405*(6788), 775–778.

Fu, L.-L., & Ferrari, R. (2008). Observing oceanic submesoscale processes from space. *Eos, Transactions American Geophysical Union*, *89*(48), 488–488.

Garrett, C. (2003). Internal tides and ocean mixing. *Science*, *301*(5641), 1858–1859.

Gaultier, L., Ubelmann, C., & Fu, L.-L. (2016). The challenge of using future swot data for oceanic field reconstruction. *Journal of Atmospheric and Oceanic Technology*, *33*(1), 119–126.

Geoffroy, G., & Nycander, J. (2022). Global mapping of the nonstationary semidiurnal internal tide using argo data. *Journal of Geophysical Research: Oceans*, e2021JC018283.

Gill, A. E., & Adrian, E. (1982). *Atmosphere-ocean dynamics* (Vol. 30). Academic press.

Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., ... Bengio, Y. (2014). Generative adversarial nets. *Advances in neural information processing systems*, *27*.

Isola, P., Zhu, J.-Y., Zhou, T., & Efros, A. A. (2017). Image-to-image translation with conditional adversarial networks. In *Proceedings of the ieee conference on computer vision and pattern recognition* (pp. 1125–1134).

Khatri, H., Sukhatme, J., Kumar, A., & Verma, M. K. (2018). Surface ocean enstrophy, kinetic energy fluxes, and spectra from satellite altimetry. *Journal of Geophysical Research: Oceans*, *123*(5), 3875–3892.

Lahaye, N., Gula, J., & Roullet, G. (2019). Sea surface signature of internal tides. *Geophysical Research Letters*, *46*(7), 3880–3890.

Lebedev, M., Vizilter, Y. V., Vygolov, O., Knyaz, V., & Rubis, A. Y. (2018). Change detection in remote sensing images using conditional adversarial networks. *International Archives of the Photogrammetry, Remote Sensing & Spatial Information Sciences*, *42*(2).

Le Guillou, F., Lahaye, N., Ubelmann, C., Metref, S., Cosme, E., Ponte, A., ... Vidard, A. (2021). Joint estimation of balanced motions and internal tides from future wide-swath altimetry. *Journal of Advances in Modeling Earth Systems*, e2021MS002613.

Lien, R.-C., & Gregg, M. (2001). Observations of turbulence in a tidal beam and across a coastal ridge. *Journal of Geophysical Research: Oceans*, *106*(C3), 4575–4591.

Metref, S., Cosme, E., Le Guillou, F., Le Sommer, J., Brankart, J.-M., & Verron, J. (2020). Wide-swath altimetric satellite data assimilation with correlated-error reduction. *Frontiers in Marine Science*, *6*, 822.

Mirza, M., & Osindero, S. (2014). Conditional generative adversarial nets. *arXiv preprint arXiv:1411.1784*.

ml4a. (2017). *Machine Learning for Artists.* Retrieved from `https://ml4a.github.io/guides/Pix2Pix/`

Morrow, R., Fu, L.-L., Ardhuin, F., Benkiran, M., Chapron, B., Cosme, E., ... others (2019). Global observations of fine-scale ocean surface topography with the surface water and ocean topography (swot) mission. *Frontiers in Marine Science*, *6*, 232.

Munk, W., & Cartwright, D. E. (1966). Tidal spectroscopy and prediction. *Philosophical Transactions of the Royal Society of London. Series A, Mathematical and Physical Sciences*, *259*(1105), 533–581.

Munk, W., & Hasselmann, K. (1964). Super-resolution of tides. *Studies on oceanography*, 339–344.

Munk, W., & Wunsch, C. (1998). Abyssal recipes ii: Energetics of tidal and wind mixing. *Deep Sea Research Part I: Oceanographic Research Papers*, *45*(12), 1977–2010.

Munk, W., Zetler, B., & Groves, G. (1965). Tidal cusps. *Geophysical Journal International*, *10*(2), 211–219.

Nelson, A. D., Arbic, B. K., Zaron, E. D., Savage, A. C., Richman, J. G., Buijsman, M. C., & Shriver, J. F. (2019). Toward realistic nonstationarity of semidiurnal baroclinic tides in a hydrodynamic model. *Journal of Geophysical Research: Oceans*, *124*(9), 6632–6642.

Park, T., Liu, M.-Y., Wang, T.-C., & Zhu, J.-Y. (2019). Semantic image synthesis with spatially-adaptive normalization. In *Proceedings of the ieee/cvf conference on computer vision and pattern recognition* (pp. 2337–2346).

Ponte, A. L., & Klein, P. (2015). Incoherent signature of internal tides on sea level in idealized numerical simulations. *Geophysical Research Letters*, *42*(5), 1520–1526.

Ponte, A. L., Klein, P., Dunphy, M., & Le Gentil, S. (2017). Low-mode internal tides

and balanced dynamics disentanglement in altimetric observations: Synergy with surface density observations. *Journal of Geophysical Research: Oceans*, *122*(3), 2143–2155.

Ponte, A. L., Le Gentil, S., & Grisouard, N. (2020). *Internal tides and geostrophic turbulence in a Boussinesq re-entrant channel.* Scholars Portal Dataverse. Retrieved from `https://doi.org/10.5683/SP2/HU58SG` doi: 10.5683/SP2/HU58SG

Ray, R. D., & Zaron, E. D. (2011). Non-stationary internal tides observed with satellite altimetry. *Geophysical Research Letters*, *38*(17).

Richman, J. G., Arbic, B. K., Shriver, J. F., Metzger, E. J., & Wallcraft, A. J. (2012). Inferring dynamics from the wavenumber spectra of an eddying global ocean model with embedded tides. *Journal of Geophysical Research: Oceans*, *117*(C12).

Rocha, C. B., Chereskin, T. K., Gille, S. T., & Menemenlis, D. (2016). Mesoscale to submesoscale wavenumber spectra in drake passage. *Journal of Physical Oceanography*, *46*(2), 601–620.

Savage, A. C., Arbic, B. K., Alford, M. H., Ansong, J. K., Farrar, J. T., Menemenlis, D., ... others (2017). Spectral decomposition of internal gravity wave sea surface height in global models. *Journal of Geophysical Research: Oceans*, *122*(10), 7803–7821.

Savva, M. A., & Vanneste, J. (2018). Scattering of internal tides by barotropic quasigeostrophic flows. *Journal of Fluid Mechanics*, *856*, 504–530.

Shakespeare, C. J., & Hogg, A. M. (2017). Spontaneous surface generation and interior amplification of internal waves in a regional-scale ocean model. *Journal of Physical Oceanography*, *47*(4), 811–826.

Shriver, J., Arbic, B. K., Richman, J., Ray, R., Metzger, E., Wallcraft, A., & Timko, P. (2012). An evaluation of the barotropic and internal tides in a high-resolution global ocean circulation model. *Journal of Geophysical Research: Oceans*, *117*(C10).

Tensorflow. (n.d.). *pix2pix: Image-to-image translation with a conditional GAN.* Retrieved from `https://www.tensorflow.org/tutorials/generative/pix2pix` (Accessed on Sep.20, 2020.)

Torres, H., Klein, P., Menemenlis, D., Qiu, B., Su, Z., Wang, J., ... Fu, L.-L. (2018). Partitioning ocean motions into balanced motions and internal gravity waves: A modeling study in anticipation of future space missions. *Journal of Geophysical Research: Oceans*, *123*(11), 8084–8105.

Torres, H., Klein, P., Siegelman, L., Qiu, B., Chen, S., Ubelmann, C., ... Fu, L.-L. (2019). Diagnosing ocean-wave-turbulence interactions from space. *Geophysical Research Letters*, *46*(15), 8933–8942.

Ward, M. L., & Dewar, W. K. (2010). Scattering of gravity waves by potential vorticity in a shallow-water fluid. *Journal of fluid mechanics*, *663*, 478–506.

Whalen, C. B., de Lavergne, C., Garabato, A. C. N., Klymak, J. M., Mackinnon, J. A., & Sheen, K. L. (2020). Internal wave-driven mixing: governing processes and consequences for climate. *Nature Reviews Earth & Environment*, *1*(11), 606–621.

Zaron, E. D. (2015). Nonstationary internal tides observed using dual-satellite altimetry. *Journal of Physical Oceanography*, *45*(9), 2239–2246.

Zaron, E. D. (2017). Mapping the nonstationary internal tide with satellite altimetry. *Journal of Geophysical Research: Oceans*, *122*(1), 539–554.

Zaron, E. D., & Egbert, G. D. (2014, 2). Time-variable refraction of the internal tide at the hawaiian ridge. *Journal of Physical Oceanography*, *44*, 538-557. Retrieved from `http://journals.ametsoc.org/doi/abs/10.1175/JPO-D-12-0238.1` doi: 10.1175/JPO-D-12-0238.1

Zaron, E. D., & Rocha, C. B. (2018). Internal gravity waves and meso/submesoscale

currents in the ocean: anticipating high-resolution observations from the swot swath altimeter mission. *Bulletin of the American Meteorological Society*, *99*(9), ES155–ES157.

Zhao, Z. (2016). Internal tide oceanic tomography. *Geophysical Research Letters*, *43*(17), 9157–9164.

Zhao, Z., Alford, M. H., Girton, J. B., Rainville, L., & Simmons, H. L. (2016). Global observations of open-ocean mode-1 m2 internal tides. *Journal of Physical Oceanography*, *46*(6), 1657–1684.

Geophysical Research Letters

Supporting Information for

A deep learning approach to extract internal tides scattered by geostrophic turbulence

Han Wang[1], Hesam Salehipour[2], Alice Nuz[1], Michael Poon[1], Aurélien L. Ponte[5], and Nicolas Grisouard[1]

[1]Department of physics, University of Toronto, Ontario, Canada

[2]Autodesk Research, Ontario, Canada

[5]Ifremer, Plouzané, France

Contents of this file

Additional Supporting Information (Files uploaded separately)

Text S1. Numerical simulations to produce snapshots in T1-5

The T1-5 simulations are based on a beta-plane centred around 45°N. The mode-1 IT is forced to the south of a baroclinically unstable jet centred in the middle section of the computational domain, and propagates northwards[1,2]. All snapshots used in the development

of TITE are publicly available, per the Data Availability Statement. Here, we summarize the relevant features of the simulations.

The baroclinically unstable jet is simulated in a zonal beta-plane channel centred at 45°N based on the primitive equation code CROCO (https://www.croco-ocean.org, v1628). Initial density profiles are different at the northern and southern ends of the domain. During a spin-up phase, the associated initial meridional density gradient undergoes geostrophic adjustment, eventually creating a zonal jet in thermal wind balance. This jet is baroclinically unstable and a zonal perturbation triggers this destabilization, resulting in low-frequency TBMs that we can reasonably describe as quasi-geostrophic[1]. Subsequently, relaxation towards unperturbed initial conditions maintains the TBMs. Statistical equilibrium is reached after *O*(100 days).

Starting at day 2000, a zonally uniform mode-1 internal tide of 12-hour period is forced within a narrow area south of the jet. Outgoing internal tides are damped in regions extending by 300 km from the southern and northern boundaries to prevent reflections back into the domain. Zonally, periodic conditions are enforced. All snapshots included in T1-5 are captured starting at day 2100. The latitudes covered by the three panels shown in Figure 1 in the main text are sufficiently away from the IT-radiating and damping regions at the southern and northern ends of the domain.

To create different levels of turbulent energy in T1 to T5, the meridional initial density gradient is modulated by changing the northern profile[2]. The TBM components are extracted online via a sliding average, replacing $\cos\left(\frac{2\pi}{T}t'\right)$ in equation (1) in the main text with a constant factor of $\frac{1}{2}$.

We compute the normalized vorticity and horizontal surface kinetic energy for the TBM components, as presented in the first three columns in Supporting Information(SI) Table S1. The TBM normalized vorticity and kinetic energy increase significantly from T1 to T5, with the normalized vorticity well bounded by 0.2.

Even though the wave amplitudes forced to the south of the turbulent jet are the same in T1-5, the IT energetics are different between simulations, due to different strengths of scattering from interactions with the TBMs. We extract the cosine IT components of surface velocities by replacing η with surface velocity components in equation (1) and compute the corresponding normalized vorticity and horizontal kinetic energy, as listed in the last two columns in SI Table S1. From T1 to T4, the kinetic energy increases. However, the kinetic energy stays about the same from T4 to T5 while getting more concentrated at smaller scales, as suggested by the increase in their respective normalized vorticities.

Moreover, the scattering from jets also makes the signals less coherent, as inspected closely by Ponte and Klein[3]. In SI Movie S3, we present snapshots of normalized vorticities of the TBMs along with the $\eta_{\cos}^{(\mathrm{sim})}$components in T1-5. There, we can see that as the TBMs become increasingly energetic from T1 to T5, the IT signals are scattered more around or to the north of the jet. As a result, the $\eta_{\cos}^{(\mathrm{sim})}$ patterns are less like plane waves and contain more small-scale features. This factor adds to the complexity of the patterns of η and $\eta_{\cos}^{(\mathrm{sim})}$.

In T1-5, the ITs are much weaker than the TBM in kinetic energy or normalized vorticity, which enables the linearized analysis conducted in past publications[2]. As a result, the internal tides are dominated by tidal wavelengths consistent with the dispersion relationship of the modal equations and the eigenvalue corresponding to the first vertical mode in the Sturm-

Liouville problem for surface fields[1,4]. The variations of density profiles in T1-5 result in variations of the tidal wavelength profiles. In the northern half of the domain, the wavelengths at higher-turbulence simulations are generally smaller than lower-turbulence simulations, as reflected in SI Fig.S1. The meridional and temporal variations of density profiles also lead to variations of the tidal wavelengths in latitude (SI Fig.S1) and time within each simulation. As density gradients are stronger in simulations at higher turbulence levels, the variation of tidal wavelengths, and hence the dominant length scales of tidal patterns, are also larger, which is another cause of the higher complexity of η and $\eta_{\cos}^{(\mathrm{sim})}$ patterns.

As mid-jet panels are centered around the baroclinic jet, the density gradients and TBMs there are on average stronger than those in the up-jet and down-jet panels. Hence, within each simulation, in the mid-jet panels, the two effects described above (scatterings of ITs and variations of tidal wavelengths) are stronger.

To sum up, the simulations correspond to a regime where TBMs, whose relative vorticities are well bounded by 1, are stronger than the ITs. The TBMs and ITs overlap significantly in spatial scales. The density profile is varied between different simulations. As a result, the TBMs and ITs become more energetic as reflected by the dynamical metrics listed in SI Table S1, and the IT wavelength profiles shift towards smaller scales, as demonstrated in Extended Data Figure 1. Enhanced scattering of ITs from TBMs causes IT incoherence, and IT patterns lose resemblance to plane waves. The density profile is varied within each simulation temporally and latitudinally, which result in corresponding variations of the IT wavelength profile. Stronger TBMs are accompanied by shaper density gradients, leading to more variations of the IT wavelengths.

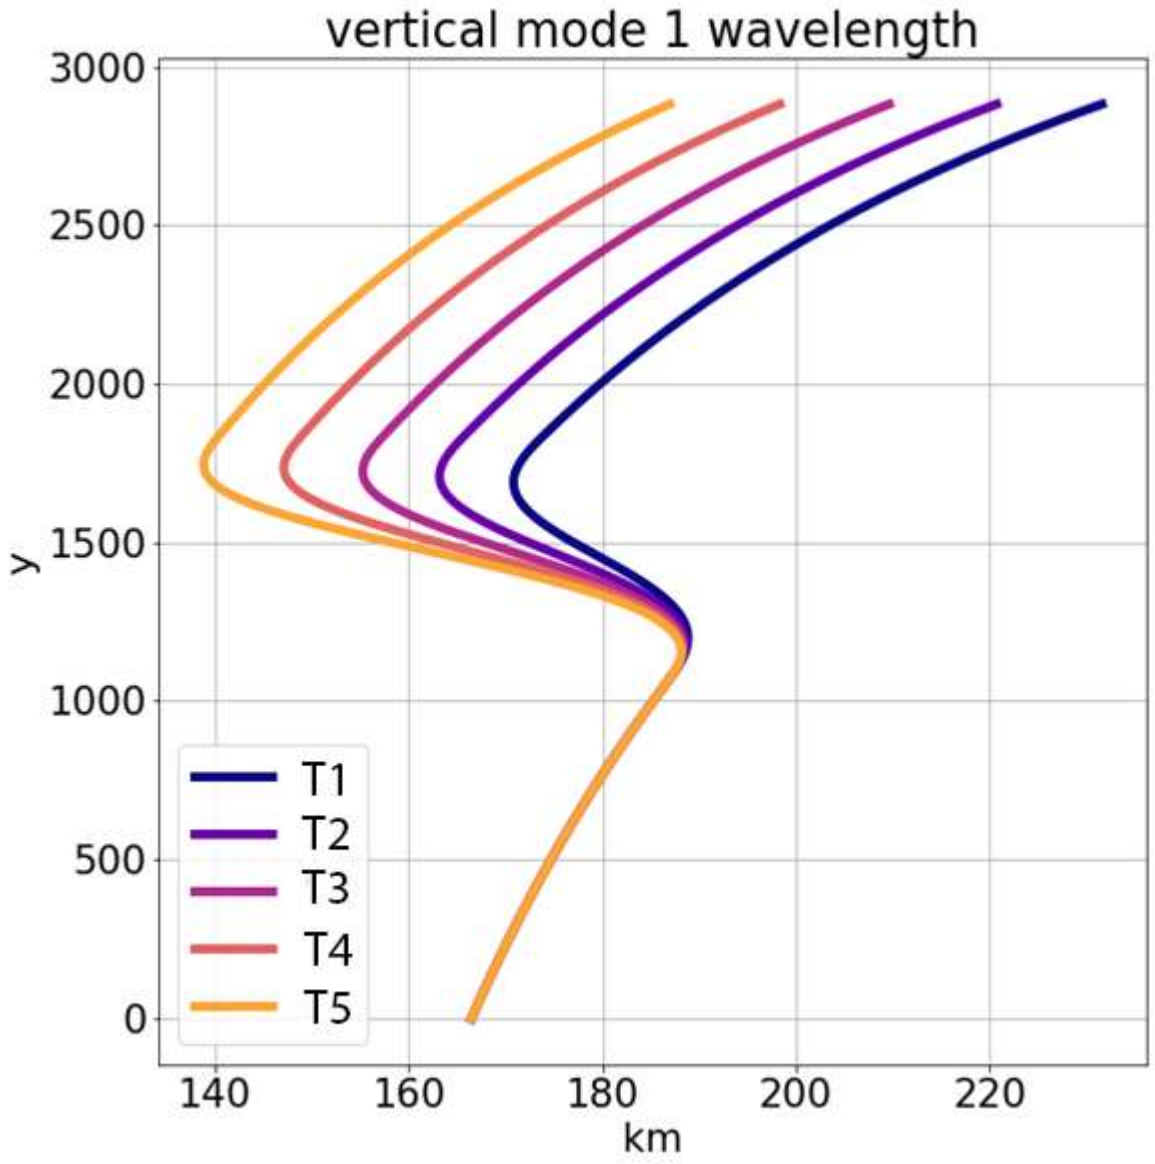

Supporting Information Fig.S1. Mode-1 wavelengths at day 1 as a function of meridional profiles. As the TBM develops, the wavelength jumps observed in the central part of the domain become smoother.

Supporting Information Table S1. Dynamical metrics of T1-T5. KE denotes "kinetic energy". The TBM normalized vorticity and KE are averaged over time, longitude, and jet width, which we define as 800 km around 45°N. The Cosine IT Normalized vorticity and KE are averaged over time and the entire simulation domain.

	TBM normalized vorticity	TBM KE (m^2/s^2)	Cosine IT normalized vorticity)	Cosine IT KE (m^2/s^2)
Simulation				
T1	0.06	0.04	2.3×10^{-3}	2.0×10^{-3}
T2	0.09	0.08	2.8×10^{-3}	2.3×10^{-3}
T3	0.12	0.12	1.6×10^{-2}	2.5×10^{-3}
T4	0.13	0.15	2.0×10^{-2}	2.6×10^{-3}
T5	0.14	0.20	2.2×10^{-2}	2.6×10^{-3}

Text S2. Detailed spectral behaviors in ET1-5 test sets

We attach in SI Fig.S2-6 the spectra and squared coherence of up-jet, mid-jet and down-jet bands in ET1-5 test sets. The squared coherence (i.e., normalized cross spectra) reflects how linearly related $\eta_{\cos}^{(\mathrm{sim})}$ and $\eta_{\cos}^{(\mathrm{gen})}$ are at different scales. It is computed based on the $\eta_{\cos}^{(\mathrm{sim})}$ and $\eta_{\cos}^{(\mathrm{gen})}$spectra following its definition listed in previous works[2]. Like the spectra, the squared coherences are computed for the up-jet, mid-jet, and down-jet bands separately in this section.

The ET1 run displays excellent spectral behaviors. In SI Fig.S2, the $\eta_{\cos}^{(\mathrm{gen})}$ spectra capture the magnitude and locations of spectral bumps of the $\eta_{\cos}^{(\mathrm{sim})}$ well, and the peaks of the squared coherence are no less than 0.8 in the up-jet, mid-jet, and down-jet bands.

In a relatively worse run (ET5) at a worse band (mid-jet band), the TITE would still outperform a simple spatial filter. This can be seen by comparing the dotted and solid yellow lines in SI Fig.S6, which correspond to the spectra of η and $\eta_{\cos}^{(\mathrm{sim})}$ signals respectively. In the η spectra, due to the strong TBMs that co-exist with ITs around the tidal wavenumbers, there is no noticeable tidal bump around the tidal wavenumber. Thus, unless one makes strong assumptions or utilizes a-priori physical knowledge, no information about the tidal wavenumber or the magnitude of tidal motions could be gained from any spatial filters applied onto η. However, TITE is still able to capture the magnitude of the spectral bumps of ITs, and the shifting between the tidal bumps of $\eta_{\cos}^{(\mathrm{sim})}$ and $\eta_{\cos}^{(\mathrm{gen})}$spectra is well less than a decade.

In all the three latitudinal bands, the squared coherences in ET1-4 decay quickly outside the tidal bumps. This is consistent with the fact that outside of the bumps, the $\eta_{\cos}^{(\mathrm{gen})}$spectra significantly mismatch the $\eta_{\cos}^{(\mathrm{sim})}$ spectra. Spurious signals in $\eta_{\cos}^{(\mathrm{gen})}$ drown out the mode-2 tidal bumps present in $\eta_{\cos}^{(\mathrm{sim})}$ and are especially prominent at large scales, sometimes causing differences between $\eta_{\cos}^{(\mathrm{gen})}$ and $\eta_{\cos}^{(\mathrm{sim})}$ by a factor of 10. As the ground truth spectra $\eta_{\cos}^{(\mathrm{sim})}$ are orders of magnitudes lower outside the tidal bumps, this decreased performance outside the bumps is not our major concern in this project. To alleviate it, we could incorporate spectral forcing in the architecture design of TITE, which is left for future work.

The squared coherences for the mid-jet band at ET5 are lower than 0.5 at all wavenumbers, suggesting the relatively poor performance of TITE in ET5 amidst the jet. Similarly, in ET2-4, the squared coherences for the mid-jet bands also peak lower than the down-jet or up-jets. An explanation of such decreased behavior around mid-jet bands is described in the main text.

The tidal bumps in the up-jet and down-jet bands are farther apart in higher turbulence runs. In ET1, the tidal bumps in the up-jet and down-jet bands almost completely overlap, while in ET5 the two bumps are shifted apart quite conspicuously. This is consistent with SI Fig.S1, which shows that the tidal wavenumbers vary more in higher turbulence runs. As discussed in the main text, the stronger variations of tidal wavelengths may be part of why the higher turbulence runs are intrinsically more challenging to TITE.

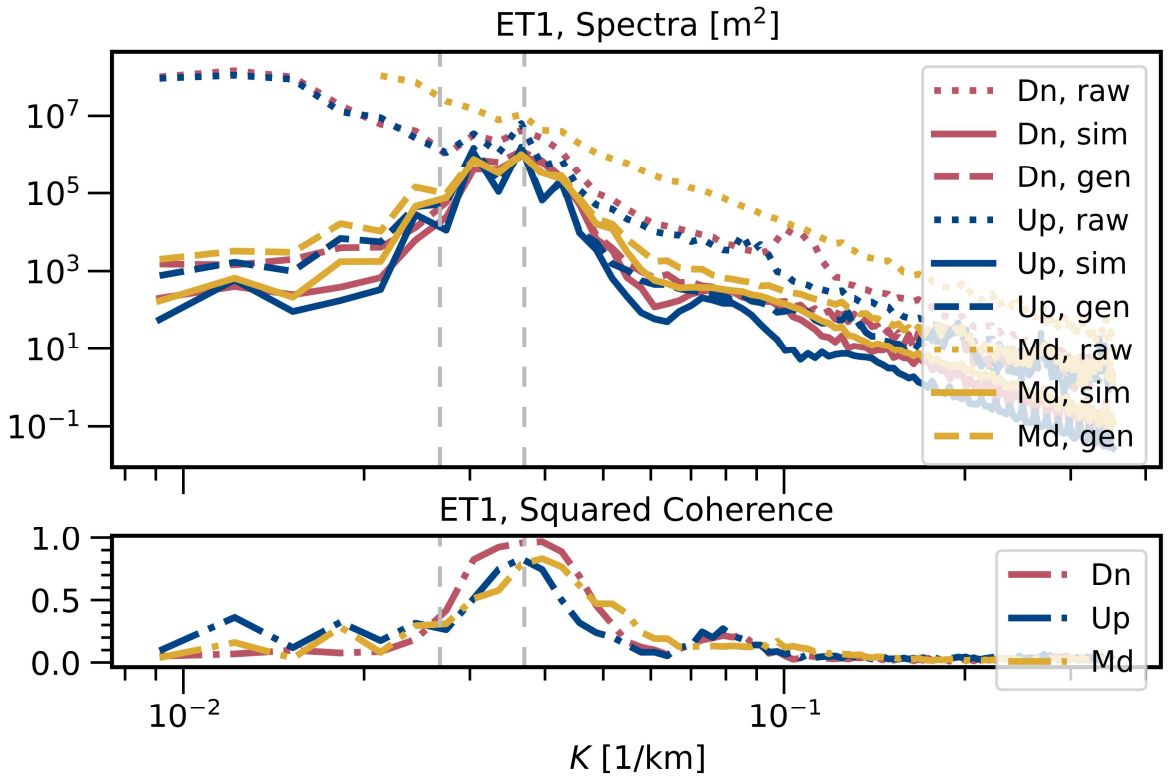

Supporting Information Fig.S2. Spectra and Coherence for the up-jet, mid-jet and down-jet bands in ET1 test set. Compared to Fig.3 in the main article, this figure presents the ET1 test set only, but adds the spectra computed for mid-jet bands (denoted by legend "Md" and coloured yellow) and the squared coherence for the three bands (lower row).

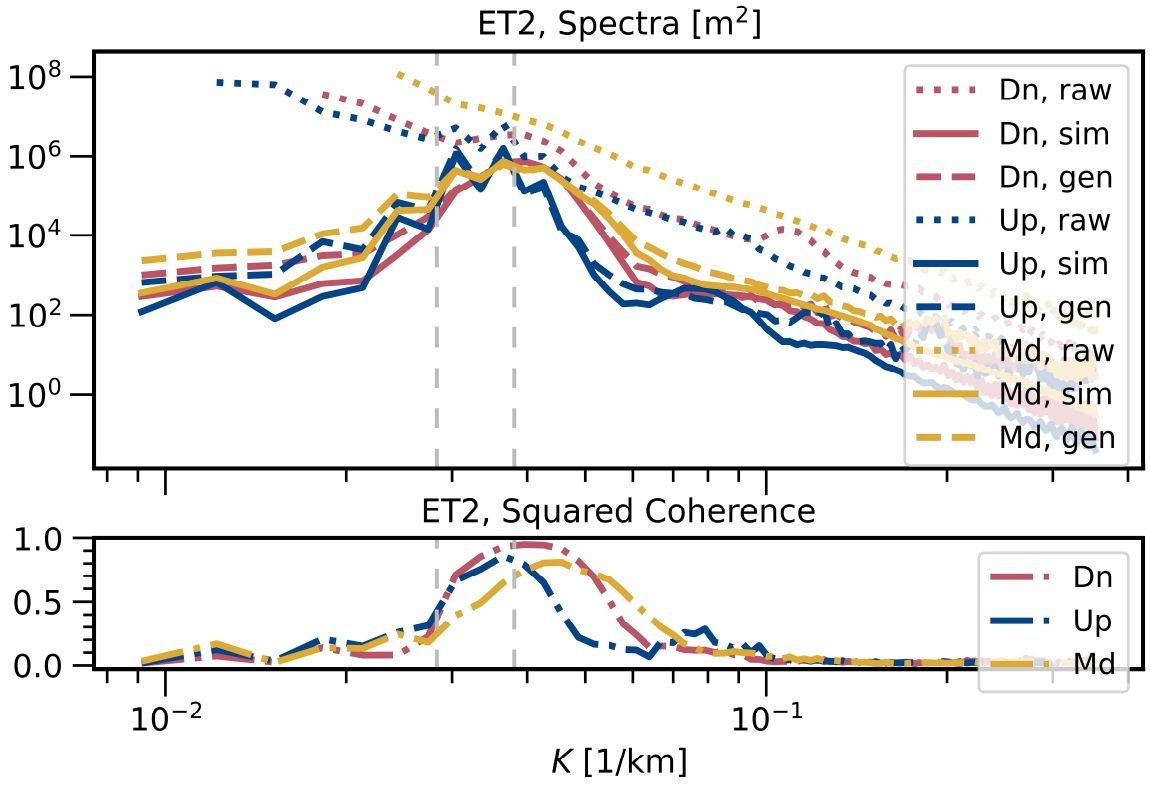

Supporting Information Fig.S3. Similar to SI Fig.S2, but for the ET2 test set.

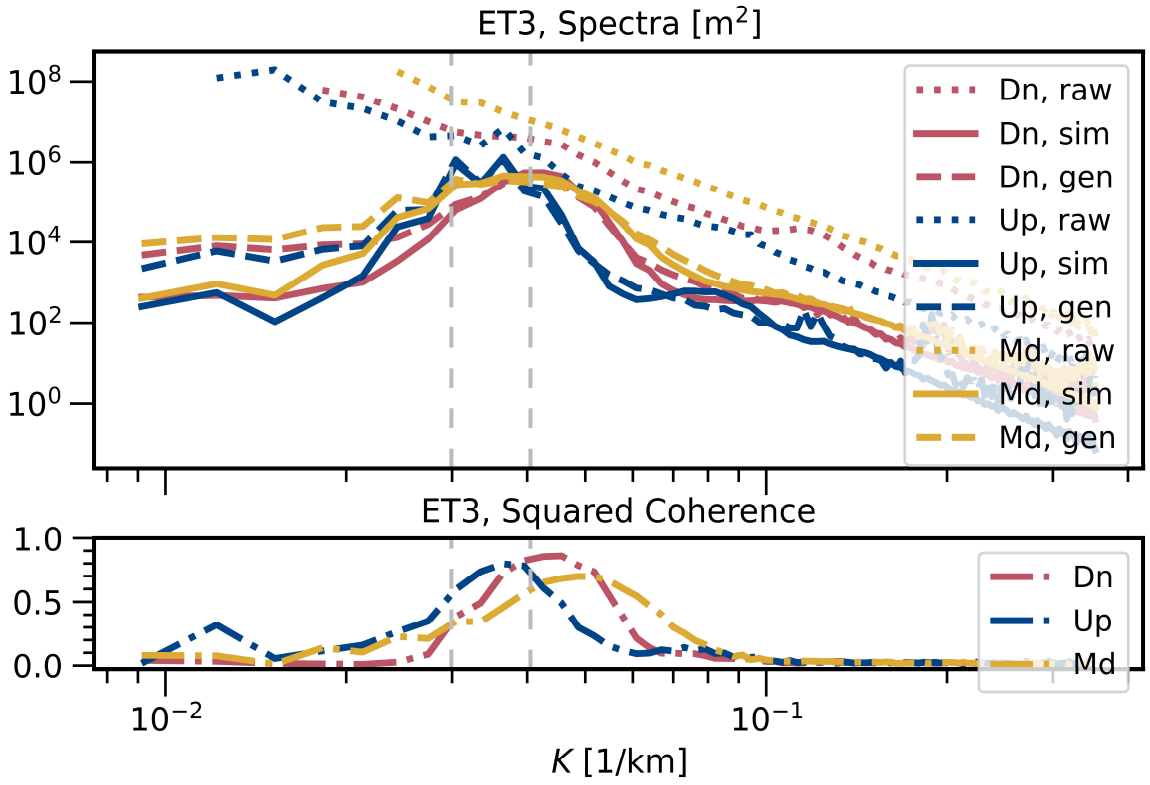

Supporting Information Fig.S4. Similar to SI Fig.S2, but for the ET3 test set.

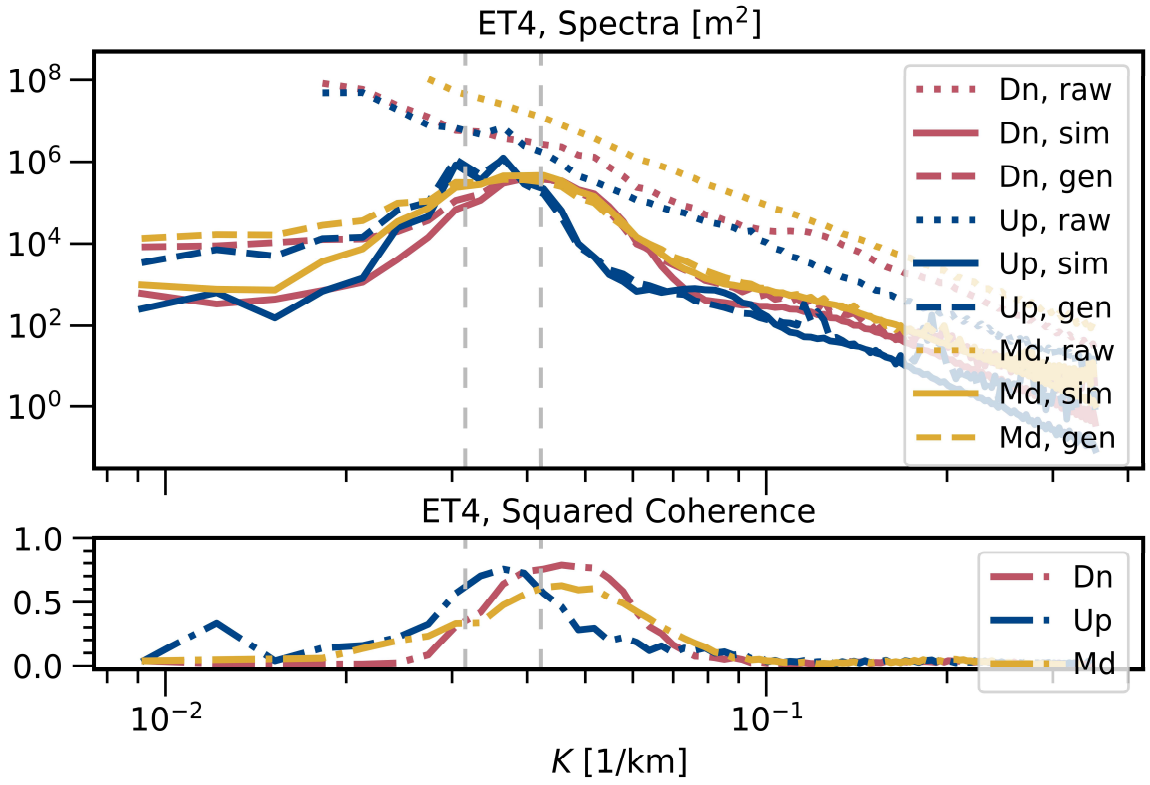

Supporting Information Fig.S5. Similar to SI Fig.S2, but for the ET4 test set.

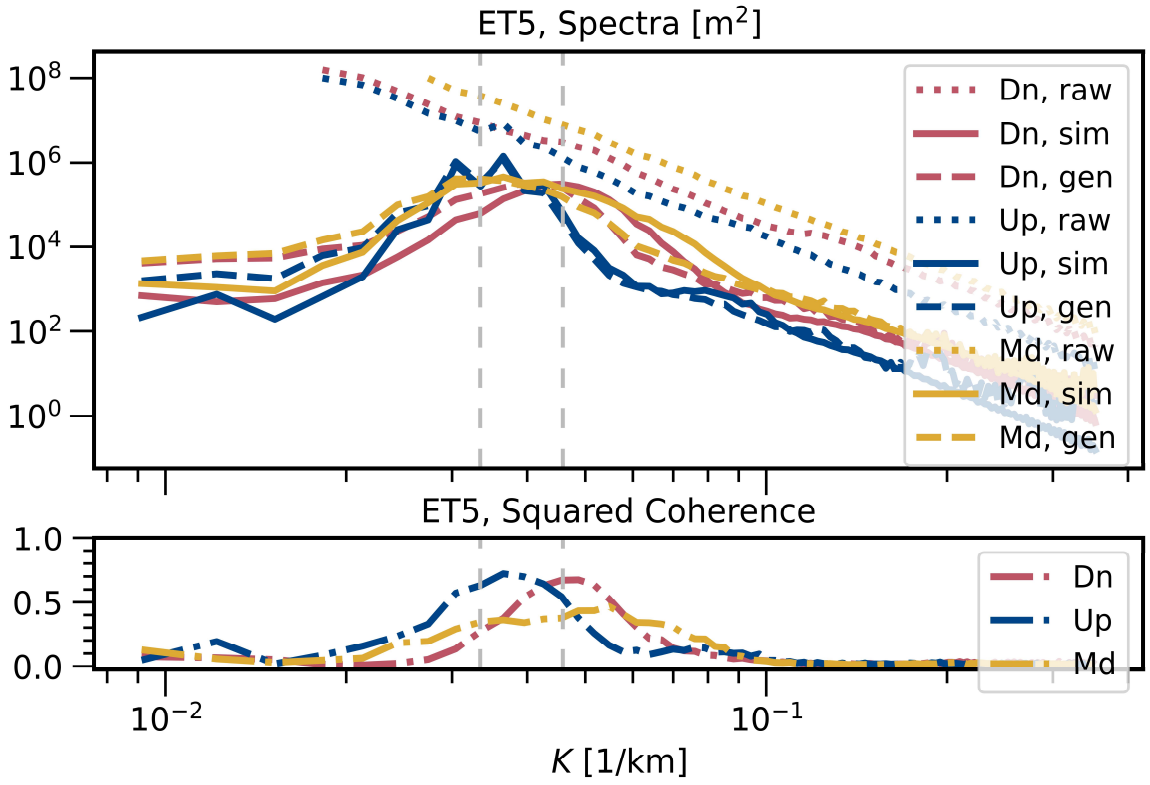

Supporting Information Fig.S6. Similar to SI Fig.S2, but for the ET5 test set.

Text S3. Statistics of correlation factors

The mean correlation values averaged over different subsets of test/validation instances are listed in SI Table S2. The histograms of the correlations in the test sets are presented in SI Fig.S7. From either the table or the histogram, the general trend of correlation to deteriorate as turbulence level gets higher, and the sharper drop from ET4 to ET5 can be observed, which are mentioned in the main text. Computations of the correlation factors are detailed in SI Text S6.

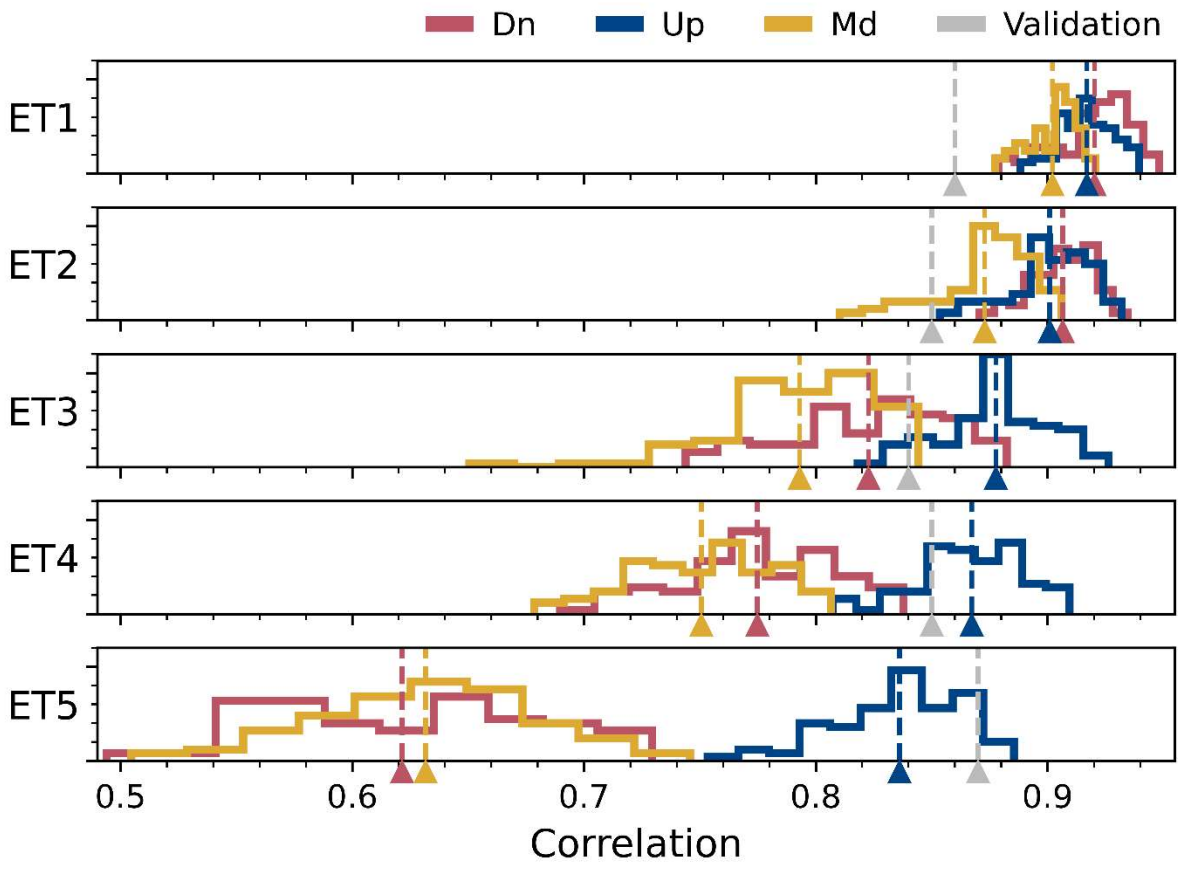

Supporting Information Fig.S7. Histogram and mean (denoted by the vertical dashed lines) of correlation factors in the test cases of ET1-5, presented for down-jet, up-jet, and mid-jet panels separately. The three panels are denoted by the colors marked in the legends, where "Dn","Up" and "Md" denote the down-jet, up-jet and mid-jet panels respectively. The mean correlations of the validation sets (averaged over all available panels) are presented in the

dashed gray vertical lines for reference. When histogram is plotted, each group is divided into 10 bins. Vertical axis group denote number counts in each bin, with axis limits fixed at 0 and 30.

Text S4. Monitoring the training and deciding the stopping epoch

The stopping criteria during a GAN training is a delicate issue, as the convergence of GAN is hard to identify due to its fleeting nature[5]. In this work, the analysis of ET1-5 in the main article are all conducted right after 600 epochs. We do not claim that it is the optimal stopping epoch for these runs but observe that there is no definite sign of model collapse around the 600th epoch.

We monitor the training behaviors from two kinds of metrics. First, we monitor the discriminator, the generator, the L1, and the total loss functions respectively as defined in the original publication[6]. Second, we monitor metrics such as correlation factor and relative error between $\eta_{\cos}^{(\mathrm{sim})}$ and $\eta_{\cos}^{(\mathrm{gen})}$ in the validation set. In SI Fig.S8, we present the evolution of discriminator loss and the correlation factor in the validation set up to the 700$^{\text{th}}$ epoch for the ET1-5 runs.

From the definition of discriminator loss in pix2pix[6], when the discriminator is effectively tossing coins at every judgement, it would have a discriminator loss of 2log(2), which is marked by the horizontal dashed line in SI Fig.S8. Observing the discriminator loss, we find that in the ET1-4 runs, the race between the discriminator and the generator appears healthily close, as the discriminator loss frequently surges above the coin-tossing line, which are then recovered back below the line in a few dozen epochs. This suggests that the discriminator and the generator are likely indeed co-evolving. Observing the correlation in the validation set of ET1-4, we see that the correlation generally stabilizes after 300th epoch with a slight tendency to increase afterwards.

The evolution of discriminator loss of the ET5 run appears less ideal. Up to the 500th epoch, the discriminator loss is always well below the coin-tossing line. In principle, this indicates a potential GAN collapse: the generator can almost never cheat the discriminator and may not be able to learn due to vanishing gradients. Between the 500th and the 700th epoch, the discriminator loss starts to surge occasionally above the coin-tossing line, which indicates that the generator may have somehow still evolved well enough to cheat the discriminator. Hence, we decide to stop at the 600th epoch, by which time the generator starts to sometimes prevail, to stay safely away from the potential collapse in earlier epochs. We note that even though the discriminator evolution is less ideal in ET5, the evolution of the correlation factor in the validation set appears to show similar behaviors as ET1-4 (bottom row, right column of SI Fig.S8), in that it stabilizes after around 300 epochs. As there are no signs of model collapsing from the evolution of the correlation, it is likely that the GAN did not collapse after all; perhaps the small bumps of the discriminator loss in the first 500 epochs in the ET5 run are sufficient to prevent vanishing gradients for the generator.

Left as future work, we can try to pace the improvement of the discriminator's performance by adding noise[5], or to use a different architecture such as the Wasserstein GAN[7] to address potentially vanishing gradients.

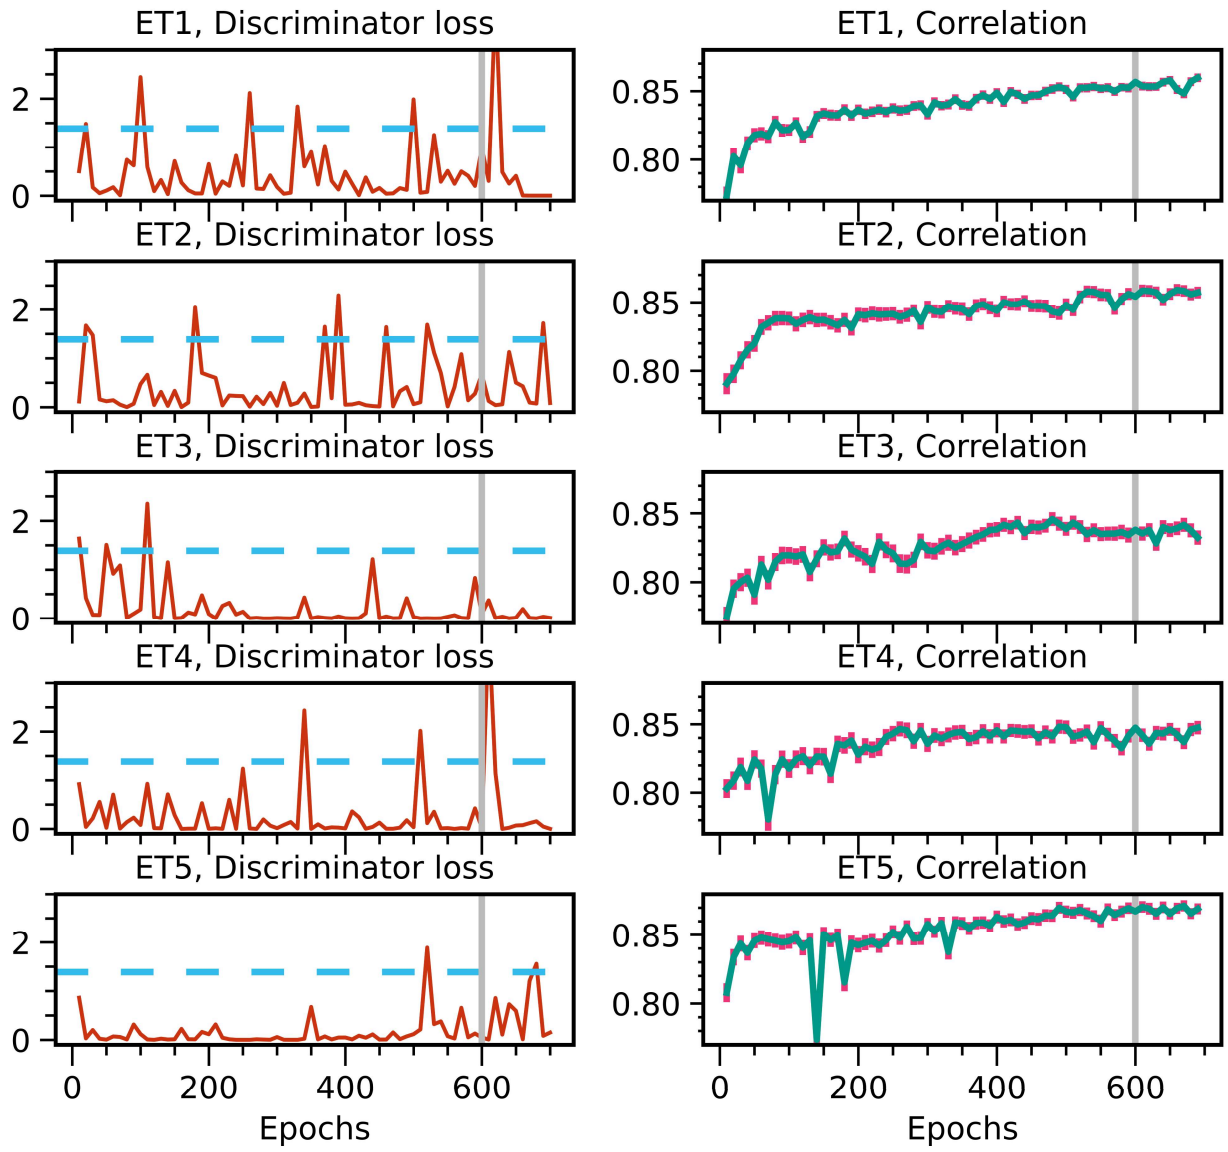

Supporting Information Fig.S8. Evolution of discriminator loss and correlation factor in validation set during the training of ET1-5. The loss and the correlation are recorded every 10 epochs, starting at the 10th epoch, and ending at the 700th epoch. Gray vertical lines mark the 600th epoch, which is the stopping epoch for the analysis in the main article. The dashed horizontal line in the left columns denote the level at which the discriminator is tossing coins. Correlation factors shown in the right columns are computed between $\eta_{\cos}^{(\mathrm{sim})}$ and $\eta_{\cos}^{(\mathrm{gen})}$ in the validation set. The line plots present the mean correlation factor of all validation instances, with error bars marking one standard deviation.

Text S5. Statistical metrics

The correlation factors and 1D spectra are computed from standard approaches. Specifically, for one panel of $\eta_{\cos}^{(\mathrm{gen})}$ and the corresponding $\eta_{\cos}^{(\mathrm{sim})}$, similar to other studies[10], we compute the correlation factor between the two arrays flattened from the two images. The mean correlation factors are averaged over all correlation factors in the data sets at interest. Take the fourth column (titled as "test set, down-jet") in Table 2 as an example. In each of the ET1-5 runs, we single out the 100 test instances belonging to down-jet panels, compute the Pearson correlation between $\eta_{\cos}^{(\mathrm{gen})}$ and $\eta_{\cos}^{(\mathrm{sim})}$ in each instance, and then average the 100 correlation factors to get the mean correlation. The maximum, minimum and standard deviation of correlation factors are computed similarly and recorded in SI 2.

Our 1D spectra are computed from 2D spectra via a numerical azimuthal averaging used in other studies [11,12]. The 2D spectra are computed over collective statistics of the down-jet, mid-jet, or up-jet panels in the test set separately. For example, in the ET5 run, the 2D spectra for the generated down-jet panels are computed from the 100 $\eta_{\cos}^{(\mathrm{gen})}$ instances from the down-jet panels in the test set. A Hanning window in the latitudinal direction is applied at each panel prior to conducting the 2D fast Fourier transforms.

In addition, we have also computed relative error of $\eta_{\cos}^{(\mathrm{gen})}$ against $\eta_{\cos}^{(\mathrm{sim})}$ for each test instance. The relative error turns out to be larger than 0.3 for each test instance in the five runs. This non-negligible relative error is consistent with the spurious large-scale signals discussed in the main article.

Text S6. Changes to the Tensorflow Tutorial code

TITE is modified from the Tensorflow Tutorial codes[13] (hereafter "tutorial codes"). Here, we detail the changes made to the tutorial codes for reproducibility. Some familiarity with the original pix2pix paper[6] from readers is assumed in the narrations to follow.

First, the $\eta_{\cos}^{(\mathrm{sim})}$ fields (ground truth) are weaker in amplitude than η (inputs) due to our simulation configurations. By trial and error, we find that this imbalance of magnitudes between inputs and outputs often destabilizes the training. To alleviate this issue, we multiply the $\eta_{\cos}^{(\mathrm{sim})}$ signals by a uniform factor of 20, after which the max value of $|\eta_{\cos}^{(\mathrm{sim})}|$ is around 78% of the max value of $|\eta|$ among all simulation snapshots we use. The other modifications we make are not essential for the training to succeed, and are rather finer improvements of training behaviours, to simplify the algorithm, or are inspired by challenges to be faced in future satellite altimetric data.

As explained in Isola et al.[6], the objective function during the training can be expressed as:

$$\arg \min_G \max_D \mathcal{L}_{cGAN}(G,D) + \lambda \mathcal{L}_{L1}(G),$$

where $\mathcal{L}_{cGAN}(G,D)$ is the classic minmax cGAN loss, and $\mathcal{L}_{L1}(G)$ is the L1 loss, which controls the impact of overall L1 error of generated images[6]. We change the parameter λ from 10^2 to 10^3, which improves the mean correlation in the validation set by around 0.09 in all the

ES1-5 runs and appears to stabilize the training. Increasing λ to 10^4 or 10^5 does not significantly change the outcomes.

As the inputs and outputs in our application are both scalar fields, we store all the panels as single-precision 2D numerical arrays rather than image-formatted files. We modified the input pipeline in the tutorial code accordingly, and the number of input and output channels is reduced from 3 (for RGB) to 1. Hence, we save some computational costs. This scalar approach is equivalent to using int32 grayscale images, and for convenience we still refer to the scalar arrays as "images" in the article. All image panels plotted in this paper are contours of the scalar fields, and the colormaps in plots are picked only for readability or aesthetic purposes. Occasionally, colours saturate in plots as an artifact from the way we define the colormaps (e.g., input fields in Fig.2), though not in our data. For normalization, we find the maximum value among all pixels in the η snapshots and divide $\{\eta, 20\,\eta_{\cos}^{(\mathrm{sim})}\}$ by this maximum value, so that all data is bounded by 1.

Prior to each epoch, training images are randomly reshuffled in time, cropped, flipped, and rotated. The random reshuffles, crops and horizontal flips are inherited from the tutorial code, whereas the random rotations and vertical flips are added by us. For random rotation, we randomly rotate each panel by 90º in either clockwise or counterclockwise directions. "Random cropping" means that we interpolate the images from a 258-by-258 to a 286-by-286 pixels grid, and within it, randomly crop a square panel of 256-by-256 pixels. All these manipulations are synchronized between the inputs η and outputs $\eta_{\cos}^{(\mathrm{gen})}$.

During random cropping, the pixel number choices of 286-by-286 and 256-by-256 are inherited from Isola et al.[6] We keep these choices for the following reasons. First, having the pixel number to be powers of 2 after cropping simplifies the downsampling steps in the generators' architecture as it helps avoid zero-paddings. Second, cropping from a 286-by-286 image to a 256-by-256 image deletes about 20% of all pixels, which is an appropriate cropping rate. The cropped images would still span over a few tidal wavelengths and thus retain the IT patterns, and yet, as the cropping causes the images to lose about 10% of the pixels in the longitudinal direction, the exact zonal periodic condition would be excluded during TITE's training, which corresponds to challenges in realistic situations

Other data augmentations (random rotations and flipping) of the training images also introduce to TITE challenges motivated by realistic situations. For example, in the simulations, ITs are forced at the southern boundary of the domain, and propagate northward. If all snapshots are upright, then during training, TITE might learn that the ITs always propagate northward, and use that knowledge during testing. But after random rotations and flipping are introduced, such information would be unavailable to TITE, which corresponds to realistic situations where one doesn't necessarily know the IT generation sites a-priori when extracting IT signals. We also experimented on TITE runs where random rotations and flipping are suppressed, and did not see any qualitative changes in TITE's performance.

Following the original nomenclature[6], our discriminator architecture can be expressed as C64-C128-C256-C512-C512-C512. The main difference between this and the architecture recommended in the original paper[6] is that at one step, our discriminator treats a whole image at once, while the original code applies a "patchGAN", which divides the image into different patches regarded independent from each other and treats each patch separately. While the patchGAN contains less convolutional layers and are less costly, one must decide on the size of

the individual patches prior to the training. We haven't investigated how to pick the patch size in our problem yet. Thus, for design simplicity, we make the patch size equal to the image size of $\eta_{\cos}^{(\mathrm{sim})}$. To investigate the impacts of this change, we have also tried using the patchGAN with the 70-by-70 patch size adopted in the tutorial code, and the mean correlation and spectral properties of $\eta_{\cos}^{(\mathrm{gen})}$ stay similar.

Additional Supporting Information

Caption for Movie S1: Performance of TITE on T1 data after trained on data from T2, T3, T4 and T5. All snapshots are re-arranged in order of time. "Input" column plots η, "Truth" column plots $\eta_{\cos}^{(\mathrm{sim})}$, "Generated" column plots $\eta_{\cos}^{(\mathrm{gen})}$, and "Difference" column plots ($\eta_{\cos}^{(\mathrm{sim})} - \eta_{\cos}^{(\mathrm{gen})}$).

Caption for Movie S2: Similar to Movie S1, but for the performance of TITE on T2 data after trained on data from T2,T3, T4 and T5.

Caption for Movie S3: Similar to Movie S1, but for the performance of TITE on T3 data after trained on data from T1,T2, T4 and T5.

Caption for Movie S4: Similar to Movie S1, but for the performance of TITE on T4 data after trained on data from T1,T2, T3 and T5.

Caption for Movie S5: Similar to Movie S1, but for the performance of TITE on T5 data after trained on data from T1,T2, T3 and T4.

Caption for Movie S6: Illustration of simulations T1-5. Five columns correspond to five simulations respectively. The upper row plots local Rossby number, defined as relative vorticities divided by Coriolis parameter. Lower row plots $\eta_{\cos}^{(\mathrm{sim})}$. The entire simulation domain is included. Snapshots are ordered by time and separated by 4*T.*

References

1. Dunphy, M., Ponte, A. L., Klein, P. & Le Gentil, S. Low-Mode Internal Tide Propagation in a Turbulent Eddy Field. *J. Phys. Oceanogr.* **47**, 649–665 (2017).
2. Ponte, A. L., Klein, P., Dunphy, M. & Le Gentil, S. Low-mode internal tides and balanced dynamics disentanglement in altimetric observations: Synergy with surface density observations: SLOW VERSUS FAST SIGNATURES ON SEA LEVEL. *J. Geophys. Res. Oceans* **122**, 2143–2155 (2017).

3. Ponte, A. L. & Klein, P. Incoherent signature of internal tides on sea level in idealized numerical simulations. *Geophysical Research Letters* **42**, 1520–1526 (2015).

4. GILL, A. E. Atmosphere-Ocean Dynamics. *Int. Geophys. Ser.* **30**, 662p (1982).

5. Arjovsky, M. & Bottou, L. Towards Principled Methods for Training Generative Adversarial Networks. *arXiv:1701.04862 [cs, stat]* (2017).

6. Isola, P., Zhu, J.-Y., Zhou, T. & Efros, A. A. Image-to-Image Translation with Conditional Adversarial Networks. *arXiv:1611.07004 [cs]* (2018).

7. Arjovsky, M., Chintala, S. & Bottou, L. Wasserstein GAN. *arXiv:1701.07875 [cs, stat]* (2017).

8. Bagrov, A. A., Iakovlev, I. A., Iliasov, A. A., Katsnelson, M. I. & Mazurenko, V. V. Multiscale structural complexity of natural patterns. *Proc Natl Acad Sci USA* **117**, 30241–30251 (2020).

9. Github, IakovlevIA. *structural-complexity,* Issue 3. https://github.com/IakovlevIA/structural-complexity/issues/3. Accessed on August 4, 2021.

10. Torres, H. S. *et al.* Diagnosing Ocean-Wave-Turbulence Interactions From Space. *Geophys. Res. Lett.* **46**, 8933–8942 (2019).

11. Torres, H. S. *et al.* Partitioning Ocean Motions Into Balanced Motions and Internal Gravity Waves: A Modeling Study in Anticipation of Future Space Missions. *Journal of Geophysical Research: Oceans* **123**, 8084–8105 (2018).

12. Savage, A. C. *et al.* Spectral decomposition of internal gravity wave sea surface height in global models. *Journal of Geophysical Research: Oceans* **122**, 7803–7821 (2017).

13. Pix2Pix | TensorFlow Core. *TensorFlow* https://www.tensorflow.org/tutorials/generative/pix2pix.