

EMERGENT PATTERNING OF MODERN ATOLL REEFS AS AN ANALOGUE FOR SELF-ORGANIZATION OF ANCIENT CARBONATES

Flora Beleznay, Gregor P. Eberli, and Sam J. Purkis

KEY FINDINGS

- Machine learning algorithms, such as convolutional autoencoders, are useful tools for assessing the patterns formed in shallow-water carbonate depositional systems.
- Specifically, the analysis of latent embeddings of dimensionally reduced data neatly demonstrates the evolution of various morphological patterns through grouping.
- Visualization of latent embeddings shows modern reef morphologies as a continuum, regardless of geographic location, suggesting shared underlying governing processes.

SIGNIFICANCE

Coherent reticulate patterns form naturally in many marine environments at scales spanning centimeters to kilometers, with one of the most studied small-scale examples being mussel beds of the Dutch tidal flats (van de Koppel et al., 2005). At larger scales, reticulate coral reef patterns have also received attention (Schlager and Purkis, 2015; Purkis et al., 2016). Yet, the origins of these patterns are often ambiguous due to similar geometries arising from external (allogenic) forcing versus internally generated (autogenic) dynamics (Burgess, 2006). By quantifying reef pattern evolution in modern atolls and assessing the role of spatial self-organization, this project establishes a mechanistic framework for the emergence of coherent patterns through biological interactions. The integration of machine learning reduces reliance on multiple spatial metrics and minimizes observational bias when comparing complex reef morphologies. These modern systems will serve as a powerful analogue to ancient carbonate depositional environments, improving interpretation of how ordered and irregular carbonate buildups form, persist, and stack in the geologic record (Purkis et al., 2015).

METHODS AND RESULTS

Reef-pattern Extraction from Satellite Imagery

To extract reef morphology, eight atolls were selected to capture a broad range of spatial patterning (*Figure 1*). High-resolution satellite imagery (3 m/pixel) and corresponding bathymetry data (10 m/pixel) were obtained from the Allen Coral Atlas (Allen Coral Atlas, 2022). This satellite imagery was derived from PlanetScope (Planet Dove) data. Because reef morphologies and environmental conditions varied substantially among sites, multiple quantitative approaches were used to distinguish reef from non-reef regions. Such methods included trained convolutional neural networks, object-based classifiers, focal statistics coupled with bathymetry data, and composite band extraction. Following reef

extraction, each site was binarized (reef/non-reef) and segmented into 2×2 km subsets (Xi et al., 2025). These binarized subsets were the input values for the autoencoder.




Figure 1: Geographic (A) and frequency (B) distributions of reticulation ranks for Goldberg's 433 'Atolls of the World' (Xi et al., 2025; Goldberg, 2016). Reefs were selected across these categories to maximize morphological diversity within the dataset.

Data Preparation

Images were initially loaded and resized to 128×128 px using nearest-neighbor interpolation, then converted to PyTorch tensors with dimensions $1 \times 128 \times 128$ px. A stratified split was performed ensuring proportional representation of each reef site across the training (70%), validation (15%), and testing (15%) datasets. To avoid overfitting, data augmentation was conducted after splitting. Augmentation ensures model reflection and rotation invariance and facilitates morphology learning independent of orientation.

Convolutional Autoencoder

The convolutional autoencoder (CAE) is composed of three essential constituents: an encoder, a latent space, and a decoder (Figure 2). The objective of a CAE is to reconstruct images from a dimensionally reduced latent representation with minimal error (Bank et al., 2023). The encoder is composed of four convolutional layers each with a kernel size of 3, padding of 1, and stride of 2, reducing the spatial dimensions by approximately half at every layer. A rectified linear unit (ReLU) activation layer follows each convolutional layer, introducing nonlinearity to allow the network to learn highly complex morphological features over simple linear transformations. The output of the encoder is a latent space of learned features from the input. The latent space is also the input value for the decoder, which simply reverses the encoder processes through upsampling, employing a transposed convolution. The final output is upsampled to the same shape as the initial input ($1 \times 128 \times 128$).




Figure 2: Architecture of the CAE used to compress binary reef morphology input. Primary components include the encoder, dimensionally reduced latent space, and decoder. PyTorch tensor values after each modification layer are included.

Model Training, Validation, and Latent Feature Extraction

Binary cross-entropy with logits loss was used as reconstruction loss to identify the best model run over 100 epochs of training. Reef (white) pixels were assigned greater weight during training to address class imbalance, coupled with an Adam optimizer to reduce reconstruction loss. After each epoch, the model is evaluated against the validation dataset before updating weights. Confusion matrices are produced at each epoch and used to calculate additional reconstruction metrics including Dice score, IoU, sensitivity, and precision. The model with the best performance according to these metrics is saved along with the raw latent space values.

Latent Space Visualization

UMAP projection was used to visualize the high-dimensional latent space using a manifold-learning approach that preserves local structure while retaining aspects of broader data organization (McInnes et al., 2018). The resulting projection reveals distinct regions and gradients of reef morphology rather than separation based solely on geographic location (Figure 3). This pattern suggests that the autoencoder is capturing shared structural characteristics among reefs across different atolls. Nearest-neighbor analysis further supports this interpretation by showing that morphologically similar reef images tend to occupy nearby regions of the latent space. Thus, modern reef morphologies appear to follow a continuum, suggesting a shared underlying governance of pattern formation regardless of environmental setting. If such relationships can be linked to underlying ecological mechanisms, this continuum may provide a modern analogue for the development of similar patterns preserved in ancient carbonate systems.




Figure 3: UMAP visualization of the autoencoder latent space for binary reef-pattern images. Each reef image is positioned according to similarity in the learned latent representation, such that samples with similar spatial morphologies occur near one another and dissimilar morphologies are more widely separated. Representative original and reconstructed images from seven reef systems are shown at right, demonstrating that the autoencoder retains the dominant spatial characteristics of each reef while simplifying some fine-scale detail.

REFERENCES

- Allen Coral Atlas, 2022. Imagery, maps and monitoring of the world's tropical reefs. Zenodo.
- Bank, D., Koenigstein, N., and Giryas, R. (2023). Autoencoders. In L. Rokach, O. Maimon, and E. Shmueli (Eds), *Machine Learning for Data Science Handbook: Data Mining and Knowledge Discovery Handbook*, 3rd ed., pp. 353-374. Springer International Publishing.
- Burgess, P.M., 2006. The signal and the noise: forward modeling of allocyclic and autocyclic processes influencing peritidal carbonate stacking patterns. *Journal of Sedimentary Research*, 76(7), pp.962-977.
- Goldberg, W.M., 2016. Atolls of the world: Revisiting the original checklist. *Atoll Research Bulletin*, 610, pp. 1-47.
- Koppel, J.V.D., Rietkerk, M., Dankers, N. and Herman, P.M., 2005. Scale-dependent feedback and regular spatial patterns in young mussel beds. *The American Naturalist*, 165(3), pp.E66-E77.
- McInnes, L., Healy, J., and Melville, J. (2018). UMAP: Uniform Manifold Approximation and Projection for Dimension Reduction. arXiv preprint arXiv:1802.03426
- Purkis, S., Casini, G., Hunt, D. and Colpaert, A., 2015. Morphometric patterns in Modern carbonate platforms can be applied to the ancient rock record: Similarities between Modern Alacranes Reef and Upper Palaeozoic platforms of the Barents Sea. *Sedimentary Geology*, 321, pp.49-69.
- Purkis, S.J., van de Koppel, J. and Burgess, P.M., 2016. Spatial self-organization in carbonate depositional environments. *SEPM Society for Sedimentary Geology Special Publication*, (106), pp.53-66
- Xi, H., Dong, X., Chirayath, V., Gleason, A.C. and Purkis, S.J., 2025. Emergent coral reef patterning via spatial self-organization. *Coral Reefs*, 44(1), pp.273-289.