

Ocean Surface Connectivity in the Arctic: Capabilities and Caveats of Community Detection in Lagrangian Flow Networks

**Key Points:**

- Community detection in the Arctic using the Infomap algorithm identifies currents as flow barriers
- Solutions found by Infomap are degenerate but ensembles of solutions are relevant
- The quality of solutions is correlated with physical quantities such as sea ice area

Supporting Information:

- Supporting Information S1

Correspondence to:

D. Reijnders,
bj.h.reijnders@uu.nl

Citation:

Reijnders, D., Jan van Leeuwen, E., & van Sebille, E. (2021). Ocean surface connectivity in the Arctic: Capabilities and caveats of community detection in Lagrangian flow networks. *Journal of Geophysical Research: Oceans*, 126, e2020JC016416. <https://doi.org/10.1029/2020JC016416>

Received 19 MAY 2020

Accepted 25 NOV 2020

Daan Reijnders¹ , Erik Jan van Leeuwen², and Erik van Sebille¹ 

¹Institute for Marine and Atmospheric Research Utrecht, Utrecht University, Utrecht, Netherlands, ²Department of Information and Computing Sciences, Utrecht University, Utrecht, Netherlands

Abstract To identify barriers to transport in a fluid domain, community detection algorithms from network science have been used to divide the domain into clusters that are sparsely connected with each other. In a previous application to the closed domain of the Mediterranean Sea, communities detected by the *Infomap* algorithm have barriers that often coincide with well-known oceanographic features. We apply this clustering method to the surface of the Arctic and subarctic oceans and thereby show that it can also be applied to open domains. First, we construct a Lagrangian flow network by simulating the exchange of Lagrangian particles between different bins in an icosahedral-hexagonal grid. Then, *Infomap* is applied to identify groups of well-connected bins. The resolved transport barriers include naturally occurring structures, such as the major currents. As expected, clusters in the Arctic are affected by seasonal and annual variations in sea-ice concentration. An important caveat of community detection algorithms is that many different divisions into clusters may qualify as good solutions. Moreover, while certain cluster boundaries lie consistently at the same location between different good solutions, other boundary locations vary significantly, making it difficult to assess the physical meaning of a single solution. We therefore consider an ensemble of solutions to find persistent boundaries, trends, and correlations with surface velocities and sea-ice cover.

Plain Language Summary To assess which surface regions of the Arctic Ocean are connected to one another, we create a division into clusters based on the exchange of virtual particles between regions. To do so, we divide the Arctic Ocean into boxes. Then, we release particles in each box, simulate their movement, and investigate the exchange between boxes. Through the computer algorithm *Infomap*, we group boxes into clusters based on their particle exchange. Knowledge about this regional connectivity is important for planning areas of conservation. We find that the boundaries between clusters often coincide with ocean currents, showing that currents provide barriers to ocean transport. We also find that the cluster division is affected by sea ice. Since the amount of sea ice in the Arctic differs across seasons and since it is decreasing over the years due to climate change, cluster divisions also change seasonally and over multiple years. Because *Infomap* returns different cluster divisions of similar quality each time it is run, we consider the common properties of multiple solutions in our analysis.

1. Introduction

Different regions of the global ocean are connected by currents and eddies. From a Lagrangian perspective, these currents and eddies facilitate the exchange of fluid parcels, that move along chaotic trajectories which change through space and time. Lagrangian ocean analysis enables studying the pathways of objects suspended in fluid (van Sebille et al., 2018), such as plastics (Hardesty et al., 2017) or the larvae of marine species (Jacobi et al., 2012; Rossi et al., 2014). With knowledge of how particles travel through different geographical areas of a fluid domain, we can investigate the connectivity of these areas.

In the context of the ocean, connectivity is a widely used term in marine ecology and marine spatial planning, relating to the larval exchange of a species between different, geographically separated subpopulations (Cowen & Sponaugle, 2009; Rossi et al., 2014). The modeling of larval dispersal has been simplified by considering larvae as passive particles (Andrello et al., 2013), sometimes also neglecting vertical effects by modeling them as buoyant particles (Rossi et al., 2014). With these simplifications, the definition of *connectivity* becomes more general, and relates to the exchange of any passive particle between geographical

© 2020. The Authors.

This is an open access article under the terms of the [Creative Commons Attribution-NonCommercial-NoDerivs License](https://creativecommons.org/licenses/by-nc-nd/4.0/), which permits use and distribution in any medium, provided the original work is properly cited, the use is non-commercial and no modifications or adaptations are made.

regions. Measures of connectivity thus become applicable to other objects that can be modeled as Lagrangian particles, such as marine debris (Maximenko et al., 2012), phytoplankton (Broekhuizen, 1999), or fluid parcels themselves. Globally, this definition of connectivity of the ocean surface has been investigated using the Lagrangian approach in the context of identifying basins of attraction (Froyland et al., 2014).

Several marine-ecological modeling studies have aimed to quantify connectivity between localized population sites by clustering sets of regions that have an internally high connectivity (Andrello et al., 2013; Jacobi et al., 2012; Thomas et al., 2014). Rossi et al. (2014) generalized this approach by not just clustering spatially separated regions, but instead considering an entire fluid domain. To do so, they describe flow in the Mediterranean Sea as a *Lagrangian flow network* and use the community detection algorithm *Infomap* (Rosvall et al., 2009) to divide the network into clusters that are sparsely connected with one another, referred to as *hydrodynamic provinces*. This approach was first presented in the context of larval exchange, since the boundaries between hydrodynamic provinces can be understood as barriers to larval transport. However, when passive buoyant particles are considered, this approach bears relevance to geophysical fluid transport in general, since boundaries between clusters can be interpreted as barriers to flow itself (Ser-Giacomi et al., 2015a). Besides clustering fluid domains, Lagrangian flow networks have also been used to assess connectivity based on particle exchange and retention between regions (Dubois et al., 2016), and to determine the most probable particle pathways between regions (Ser-Giacomi et al., 2015b).

In recent years, much work has aimed to detect Lagrangian coherent structures (LCSs), which are delineated by material lines that are linearly stable or unstable for longer times than surrounding regions (Haller & Yuan, 2000). LCS detection methods are varied in nature and rely on different underlying definitions of what constitutes an LCS (see Hadjighasem et al. [2017] for a comparison). Some methods provide diagnostic scalar fields that highlight coherent structures, while others yield partitions of a domain in which each cluster is maximally coherent (or equivalently, the exchange between clusters is minimized). Hydrodynamic provinces are not only characterized by small fluid exchange across their boundaries, but also by high-internal mixing (Rossi et al., 2014; Ser-Giacomi et al., 2015a), which bears relevance to marine ecology in particular and implies a high reachability of different locations within hydrodynamic provinces over time. The high-internal mixing is an important distinction between LCSs and the hydrodynamic provinces investigated here. The two may bear some relations: attracting and repelling material lines govern which regions of a domain are reachable for particles, thus influencing connectivity and, in turn, the topology of hydrodynamic provinces. A full comparison, however, is beyond the scope of this research.

We aim to investigate surface connectivity in the Arctic and subarctic oceans through the identification of hydrodynamic provinces using the Lagrangian flow network approach. We confine ourselves to the surface, which is relevant for buoyant particles, such as larvae and marine debris. The Arctic provides an interesting domain of study for several reasons. First, to our knowledge, no previous clustering studies have been carried out in the Arctic domain. Moreover, the *Infomap* algorithm has thus far only been applied to closed-domain Lagrangian flow networks (Rossi et al., 2014; Ser-Giacomi et al., 2015a), while the Arctic and subarctic Oceans comprise a domain that is, open at the southern boundary. We thus test whether this approach is successful at identifying meaningful communities in open domains.

Second, the Arctic Ocean experiences strong seasonal variations in the strength and location of ocean currents. These currents may hinder transport between different regions, and it is insightful to see whether they thus act as boundaries between hydrodynamic provinces. Currents that shape the circulation in the ice-free regions of the Arctic and that we expect to affect clustering patterns include the Norwegian Current, the West Spitsbergen Current, the East Icelandic Current and the Irminger Current. Within regions containing sea ice, the Beaufort Gyre dominates the surface flow pattern. An overview of important currents and oceanic structures in the Arctic is presented in supporting information Text S1. The Arctic Ocean is also subject to seasonal variations in the sea ice extent, with sea-ice affecting surface flow as well (Goosse & Fichefet, 1999). These variations influence the location of barriers to flow. We compare connectivity between different seasons and years, in order to see which physical mechanisms are governing barriers to transport.

Lastly, the average Arctic sea ice extent has been decreasing over the past decades and is very likely to decrease in the future (Comiso, 2012; Vaughan et al., 2013). Since this decrease will cause a potentially

irreversible shift into a new climatic state (Polyakov et al., 2013), it is insightful to see whether these developments are reflected in the topology of hydrodynamic provinces.

We also aim to describe important considerations when using this approach and to raise caveats that have not been previously discussed. This includes a discussion of the physical interpretation of communities found by the community detection algorithm *Infomap* (Rosvall et al., 2009), which may also be relevant to other community detection methods in flow networks. Moreover, community detection algorithms in complex networks have been shown to be sensitive to degenerate solutions, meaning that many good solutions may exist, while their topology may significantly differ (Calatayud et al., 2019; Good et al., 2010). We therefore, assess which structures are persistently found across different solutions.

2. Theory

2.1. Constructing Lagrangian Flow Networks

By mapping flow onto a network, the dynamics of the fluid system are captured by the topology of the network (Molkenthin et al., 2015). The explicit characterization of a fluid as a network constructed from Lagrangian trajectories was first introduced by Rossi et al. (2014) and later described from a more technical perspective by Ser-Giacomi et al. (2015a). This method constructs *Lagrangian flow networks* based on particle or mass exchange between boxes in a discretized domain. Other methods define network representations based on distances between Lagrangian particle trajectories (Hadjighasem et al., 2016; Wichmann et al., 2020), or close encounters of such trajectories (Padberg-Gehle & Schneide, 2017). A discussion of different methods to construct network representations of flow systems is found in Donner et al. (2019).

The Lagrangian flow network characterization enables us to analyze fluid dynamics using the vast toolbox of network science (see Newman [2010] for an overview). A frequently recurring problem in network science is the division of a network into communities of nodes that are well connected among each other, with only sparse connections between distinct communities (Newman & Girvan, 2004). For comparisons between the many approaches that have been proposed for solving this problem, see Danon et al. (2005) and Fortunato (2010). In the context of Lagrangian flow networks, ideally such a network division yields barriers to fluid transport, with fluid being unlikely to cross these barriers. Simultaneously, high connectivity within a network community should correspond to the fluid within one community being well-mixed.

A network representation of a system comprises a graph $G = (V, E)$, consisting of a set of *nodes*, V , and a set of *edges*, E , where an edge $(i, j) \in E$ forms a connection between nodes $i, j \in V$. Additionally, these edges may be directed, so that an edge from node i to node j is distinct from an edge from j to i . Moreover, edges may take on weights w_{ij} , which can correspond to the importance of a connection or the relative volume of the flow. In our practical application, V and E are finite sets.

When mapping fluid flow as a network, the fluid domain needs to be discretized in order to represent the continuous flow by the finite sets V and E . We can divide the flow domain into a set of N_B bins, $B = \{B_i, i = 1, \dots, N_B\}$, and consider the flow between different bins. These bins then form the nodes of the network, while the flow between bins is captured by the edges between nodes. Since the flow between bins is directional and can differ in magnitude, edges should be weighted and directed.

The flow between bins depends on the initial state of the fluid at time t_0 and the time interval τ in which the flow is considered. In Lagrangian flow networks, we establish a connection between node i and node j if there is exchange of fluid from the corresponding bin B_i to bin B_j in the time interval $[t_0, t_0 + \tau]$. The weight of edge (i, j) is taken proportional to the amount of fluid that is, transported from B_i to B_j .

From a Lagrangian perspective, fluid transport can be approximated from the initial and final positions in the trajectories of ideal fluid particles. These trajectories can be determined through integration of the equations of motion of particles. Final particle positions $\mathbf{X}(t_0 + \tau)$ are then given by

$$\mathbf{X}(t_0 + \tau) = \mathbf{X}(t_0) + \int_{t_0}^{t_0 + \tau} \mathbf{v}(\mathbf{x}(t), t) dt, \quad (1)$$

where $\mathbf{v}(\mathbf{x}, t)$ is the time-dependent Eulerian velocity field (van Sebille et al., 2018). Then, the right-hand side of Equation 1 defines the flow map $\Phi_{t_0}^\tau$, which maps the initial location x of a fluid particle to its final location.

Given $m(B_i)$ Lagrangian particles initially being distributed in bin $B_i \in B$, we can approximate the flow probability between bins by considering the fraction of particles traveling from bin B_i to bin B_j in time window $[t_0, t_0 + \tau]$ by

$$\mathbf{P}(t_0, \tau)_{ij} = \frac{\#\{x : x \in B_i \text{ and } \Phi_{t_0}^\tau(x) \in B_j\}}{m(B_i)}, \quad (2)$$

which allows us to construct a transition matrix $\mathbf{P}(t_0, \tau)$ (Froyland et al., 2014; Rossi et al., 2014; Ser-Giacomi et al., 2015a). Therefore, $\mathbf{P}(t_0, \tau)$ defines the approximation of our flow as a Markov chain. As long as fluid parcels, or equivalently, particle trajectories are conserved, $\mathbf{P}(t_0, \tau)$ is row-stochastic, such that for each bin B_i , we have

$$\sum_{j=1}^{N_B} \mathbf{P}(t_0, \tau)_{ij} = 1 \quad (3)$$

and each element is nonnegative. Note that the definition of our transition matrix (2) only considers the initial and final location of the particles, and thus contains no information about fluid exchange between bins at intermediate times.

The transition matrix $\mathbf{P}(t_0, \tau)$ can be used as an adjacency matrix to generate a graph $G_{t_0}^\tau$, which is our network representation of the flow. Each row and column index corresponds to a node, and the weight of an edge $w(i, j)$ is given by the entry $\mathbf{P}(t_0, \tau)_{ij}$. In the network representation of the fluid, edge weights thus correspond to the probability that a particle travels between bins and the row-stochastic property of $\mathbf{P}(t_0, \tau)$ ensures that the sum of the weights of outgoing edges for any given node is 1.

2.2. Community Detection and Infomap

Hydrodynamic provinces should be characterized by little fluid exchange between provinces, while their interior should be well-mixed, meaning that fluid from one location in the hydrodynamic province is exchanged evenly to other locations in the province. These two criteria can be translated into concepts from network science, where they correspond to the goal of community detection (Ser-Giacomi et al., 2015a). Respectively, different communities should be sparsely connected, while communities themselves correspond to well-connected regions in the graph.

There is no single definition of what constitutes a community. Some networks exhibit clear community structures that can be identified visually, while some networks exhibit no clear community structure at all. Even when there exist spatial differences in connectedness that hint at the presence of communities, it often is not apparent where community boundaries should be drawn exactly. Moreover, it is important to determine a spatial scale at which the investigation of communities can lead to meaningful results, since communities may exhibit a nested, multilevel structure (Kheirkhazadeh et al., 2016). A definition of what constitutes a community implicitly depends on the detection strategy used, causing different algorithms to detect community structures of different nature (Rosvall et al., 2017).

One popular family of methods for detecting communities is modularity maximization (Newman, 2010; Newman & Girvan, 2004). This strategy relies on the comparison of a given network with a random network or another null model, in order to determine which regions of the network exhibit more connections than would be expected in a random network. While modularity maximization is popular, it requires null models which carry no obvious meaning with respect to flow networks, thus lacking a physical interpretation (Ser-Giacomi et al., 2015a). In addition, these methods suffer from a *resolution limit*, preventing us from detecting communities smaller than a specific scale that is, determined by the size of the total network (Fortunato & Barthelemy, 2007).

Instead, Rossi et al. (2014) and Ser-Giacomi et al. (2015a) propose using the *Infomap* community detection method (Rosvall et al., 2009) for detecting hydrodynamic provinces. *Infomap* is a community detection algorithm that takes edge directionality and weights into account and its solutions are less likely to be impacted by a resolution limit than other methods (Kawamoto & Rosvall, 2015). In addition, it can find communities that may differ in size. *Infomap* also allows the study of community structures at different scales, either through identifying nested communities (Edler et al., 2017), or through a tuning parameter that affects community sizes (Kheirkhahzadeh et al., 2016). It was first introduced by Rosvall and Bergstrom (2008) and was expanded on and presented as a software package by Rosvall et al. (2009).

Infomap uses an information-theoretic approach to identify communities based on how flow within a network is constrained by its topology. *Infomap* simulates flow by considering the movement of random walkers on a network. Specifically, it aims to partition the network into communities such that it minimizes the average length of encoded trajectories of random walkers, which traverse the network with probabilities corresponding to the local importance of edges. For our flow network, edge weights directly correspond to the transition probabilities of Lagrangian particles. The transition probabilities of Lagrangian particles will thus be used by *Infomap* to drive the movement of the random walkers. We will briefly explain the basic principles and intuitions behind *Infomap* and subsequent extensions (Kheirkhahzadeh et al., 2016; Lambiotte & Rosvall, 2012) that are important for understanding our results. A more detailed account can be found in the supporting information (see Text S2).

Infomap capitalizes on the fact that trajectories of random walkers can be encoded into strings of bits by using Huffman codes (Huffman, 1952), an optimally efficient encoding method. A particle's trajectory can be described as a sequence of visited nodes in the flow network. Each node is assigned a unique string of bits, called a *codeword*. Huffman codes are optimally efficient by assigning short codewords to nodes that have a high ergodic visiting frequency, and longer codewords to nodes that are visited less often. An added constraint is that no codeword can be the prefix of another, limiting the set of possible codewords. Through this method, we can communicate a node-to-node trajectory of a random walker using a concatenated sequence of codes, which we refer to as a path description. By assigning shorter codewords to frequently visited nodes, the average length of a path description is minimized. The length of codewords will grow in size as the number of nodes in the network increases, as more bits are needed to describe a node using a prefix-free codeword. This in turn leads to longer path descriptions.

Rather than describing a network only in terms of nodes and edges, we can consider the nodes in a network to be divided into communities. When a network is made up of communities that are characterized by few edges between communities and many edges within a community, a random walker is likely to spend a long time within a community before moving to another. This allows the construction of a two-level encoding system, using a separate encoding for the events of entering each community and for the movement of a random walker within a community. For the latter, each community is given its proper encoding, allowing codewords to be reused between different communities. Since codewords are reused, they can be shorter. It is important that random walkers are unlikely to switch between communities, since that would require sending an extra codeword recording this switch, thus increasing the length of a path description. Finding a good division into communities will thus yield shorter path descriptions. This implies a minimization strategy for finding good communities, namely to find a community partition that minimizes the average length of codewords used in a random walk.

Rather than explicitly determining the most efficient encoding for a given community division, *Infomap* instead takes advantage of concepts from information theory to find the theoretical lower bound of the average codeword length (codelength). Given a partition \mathbf{P} that divides the n nodes in V into c communities $\alpha = 1, 2, \dots, c$, this lower bound is denoted by $L(\mathbf{P})$. To find an expression for this lower bound, Shannon's source coding theorem is used (Shannon, 1948), which implies that the average length of a codeword is bounded from below by the entropy of the random variable X , the n states of which are described by n distinct codewords. With p_i denoting the frequency of occurrence of a state, the Shannon entropy is then given by

$$H(X) = -\sum_{i=1}^n p_i \log_2(p_i). \quad (4)$$

The information-theoretical lower bound on the average length of a codeword describing a step of the random walk is then found through a weighted average of the entropy associated to the length of the codewords describing entering a new community and the entropies corresponding to the codewords describing steps within each community. This is captured in the map equation:

$$L(\mathbf{P}) = q_{\curvearrowright} H(\mathcal{Q}) + \sum_c^{\alpha=1} p_c^{\alpha} H(\mathcal{P}^{\alpha}). \quad (5)$$

Here $H(\mathcal{Q})$ is the frequency-weighted average codelength corresponding to switching between communities, while $H(\mathcal{P}^{\alpha})$ is the frequency-weighted average codelength describing steps within a community α . These entropy terms are respectively weighted by the probability that a random walker exits a community, q_{\curvearrowright} , and the probability of using the codes corresponding to steps in community α , denoted by p_c^{α} . Instead of evaluating these terms by simulating the movement of random walkers on the network, they can be calculated from the transition matrix $\mathbf{P}(t_0, \tau)_{ij}$. This is explained in detail in the supporting information (Text S2).

We can tune the typical size at which *Infomap* detects community structures in a network. This allows us to examine connectivity at a scale that is, useful for investigating oceanographic structures. This is done through the Markov-time parameter t_m , which governs the amount of steps a random walker is allowed to take before its position is recorded into the path description (Schaub et al., 2012a; Schaub 2012b). Generally, a higher value of t_m means the random walker can travel more, thus making it more likely to transition to different areas of the network and thus increasing the average size of communities in order to keep the average codelength short. Note that the typical size of a community is also determined by the integration time, with longer integration times yielding more particle dispersal, thus connecting more areas of the network. Increasing the Markov-time, in contrast, tunes the spatial scale of connectivity without considering the flow itself at longer time scales. An account of how t_m relates to the map Equation 5 is found in the supporting information (Text S3).

Infomap uses a stochastic and recursive heuristic algorithm to minimize the map equation. Its core algorithm roughly follows the following steps. Initially, each node is assigned its own community. Then, in random order, each node is moved to the neighboring community that would reduce L the most (see Equation 5), unless no move reduces L , in which case the node remains in its original community. It then applies this iteration recursively until no move results in a reduction of L . After that, *Infomap* is recursively applied on the resulting partition, now using the communities as nodes, until L can be no longer reduced. This causes *Infomap* to find local minima of L .

The use of a heuristic and stochastic approach has important implications. Most importantly, due to its stochastic nature, different passes of *InfomSap* can yield different locally optimal solutions. The algorithm may be run multiple times such that the partition \mathbf{P} that yields the lowest value of $L(\mathbf{P})$ can be picked as a final solution. However, many degenerate solutions may exist, which all have similar values of L while they may exhibit considerable topological differences (Calatayud et al., 2019; Good et al., 2010). The transition matrices used by *Infomap* are by themselves already approximations of the real surface flow. This means that when one solution \mathbf{P}_a has a slightly lower value of L than another solution \mathbf{P}_b while exhibiting a significantly different topology, there is no reason to assume that solution \mathbf{P}_a carries more physical meaning than \mathbf{P}_b . Investigating the structure of only one solution might be misleading, especially when a community structure is weak (Calatayud et al., 2019). While some communities are consistently found across different good solutions, others may not. Different solutions can be merged to find a consensus solution (Lancichinetti & Fortunato, 2012; Strehl & Ghosh, 2003), but in doing so, information on which community boundaries are weak may be lost. Instead, it is insightful to compare multiple solutions to see on which structures solutions agree and to figure out in which regions of the network the community structures are weaker (Calatayud et al., 2019).

2.3. Quality of Hydrodynamic Provinces

Infomap's sole criterion for finding a good partition \mathbf{P} is minimizing $L(\mathbf{P})$, which carries no obvious physical meaning. We can assess the quality of solutions found by *Infomap* through two other measures, relating to the goals of community detection. First, the ratio between Lagrangian particles leaving and staying in a hydrodynamic province within time τ should be low. This criterion interprets hydrodynamic provinces as

almost-invariant areas of fluid, such that flow within a region A is nearly mapped onto itself after time τ : $\Phi_0^\tau(A) \approx A$ (Ser-Giacomi et al., 2015a). Second, hydrodynamic provinces should have strong internal mixing, making sure that different areas of each hydrodynamic province exchange fluid.

Ser-Giacomi et al. (2015a) propose two quality parameters to assess the extent to which these criteria are met. The first criterion is assessed through the *coherence ratio*, $\rho_0^\tau(\alpha)$, measuring the ratio between particles that leave and stay within a community α within time step τ .

$$\rho_0^\tau(\alpha) = \frac{\sum_{i,j \in \alpha} m(B_i) \mathbf{P}(t_0, \tau)_{ij}}{\sum_{i \in \alpha} m(B_i)}. \quad (6)$$

For a partition \mathbf{P} that divides the domain into c communities $\alpha = 1, \dots, c$, the global coherence ratio is the average of the coherence ratio of each community:

$$\rho_0^\tau(\mathbf{P}) = \frac{1}{N_B} \sum_c^{\alpha=1} \#\{B_i \mid i \in \alpha\} \rho_0^\tau(\alpha). \quad (7)$$

Unlike in Ser-Giacomi et al. (2015a), here the global coherence ratio is weighted by the amount of bins in a community, such that we minimize the effect of small communities produced by noise in the data. The coherence ratio is determined only through $\mathbf{P}(t_0, \tau)_{ij}$, which is constructed from only the initial and final particle locations, meaning that particles may temporarily leave a community within the interval $[t_0, t_0 + \tau]$.

The second criterion is assessed using a measure of mixing proposed by Ser-Giacomi et al. (2015a). The *mixing parameter* indicates how strongly fluid within a community is mixed. To do so, only flow occurring within a community α is considered, which we can represent through a reduced transition matrix

$$\mathbf{R}(t_0, \tau \mid \alpha)_{ij} = \frac{\mathbf{P}(t_0, \tau)_{ij}}{\sum_{k \in \alpha} \mathbf{P}(t_0, \tau)_{ik}}, \quad i, j \in \alpha. \quad (8)$$

The mixing parameter for a community $M_{i_0}^\tau(\alpha)$ is given by the normalized sum of the Shannon entropy associated to the transition probabilities between each pair of bins:

$$M_{i_0}^\tau(\alpha) = \frac{-\sum_{i,j \in \alpha} \mathbf{R}(t_0, \tau \mid \alpha)_{ij} \log \mathbf{R}(t_0, \tau \mid \alpha)_{ij}}{Q_\alpha \log Q_\alpha}, \quad (9)$$

with $Q_\alpha = \#\{B_i \mid i \in \alpha\}$. The mixing parameter reaches its maximum value of one when particles within a bin B_i , $i \in \alpha$ are dispersed uniformly to all other boxes in α ($\mathbf{R}_{ij} = \frac{1}{Q_\alpha} \forall i, j \in \alpha$). The global mixing parameter is then the weighted average of the mixing parameter

$$M_{i_0}^\tau(\mathbf{P}) = \frac{1}{N_B} \sum_c^{\alpha=1} Q_\alpha M_{i_0}^\tau(\alpha). \quad (10)$$

For practical applications such as investigating barriers to transport, a third, qualitative, criterion can be added, namely that the communities found by *Infomap* take on spatial scales that are useful for identifying these barriers. The Markov-time parameter t_m in *Infomap* allows us to change the spatial scale for investigation, and may thus be used to fulfill this criterion. This criterion is not considered by Ser-Giacomi et al. (2015a). It cannot be stated a priori what value of t_m should be used. Instead, it is a tuning parameter that can be adjusted to find structures at an appropriate spatial scale. It can be tuned after inspecting the spatial scale of communities found by *Infomap* for the default value $t_m = 1.0$.

3. Materials and Methods

Our methods for finding hydrodynamic provinces from a Lagrangian flow network closely follow the approach by Rossi et al. (2014) and Ser-Giacomi et al. (2015a) as described in the previous section. Machine configuration and running times are found in the supporting information (Text S4).

3.1. Hydrodynamical Data

To describe the hydrodynamics in the Arctic we use the Global Ocean Physical Reanalysis product (GLOBAL_REANALYSIS_PHY_001_030) (Fernandez et al., 2018), made available by the Copernicus Marine Environment Monitoring Service (CMEMS). This product provides reanalysis data for the global ocean at a resolution of $1/12^\circ$, corresponding to a latitudinal length of 9.3 km per grid cell. However, the first baroclinic Rossby radius of deformation, which is the natural scale of baroclinic boundary currents, eddies and fronts takes values between 1 and 16 km in the Arctic Ocean, sometimes even assuming values below 1 km in shallow seas like the Barents Sea (Nurser & Bacon, 2014). Therefore, eddies, fronts and boundary currents of these scales are not resolved in certain regions of the Arctic. Instead, we resolve larger-scale or aggregate structures larger than our grid size.

While hydrodynamical data is provided for 50 vertical levels, we only use the uppermost layer at 0.49 m depth to describe dynamics at the surface, which is relevant to buoyant particle transport. The data set contains daily mean fields over the period 1993–2018, for which ocean surface altimetry data and satellite sea ice data are available. This enables us to investigate how seasonality and annual trends influence the topology, coherence and mixing of hydrodynamic provinces. Moreover, the temporal resolution and extent allows us to investigate the persistence of features over time.

The reanalysis product is constructed by assimilating model output from the NEMO 3.1 ocean model (NEMO System Team, n.d.) and the LIM2 EVP sea ice model (Goosse & Fichefet, 1999) with observational data. While values in NEMO and LIM2 are computed on a tripolar Arakawa C-grid, final fields are interpolated on a regular Arakawa A-grid. This is a reason for caution, since a resolution increase on the interpolated A-grid as we move poleward does not correspond to smaller scale structures being resolved any better. In fact, closer to the pole, many grid cell values may be interpolated from just a few cells on the C-grid.

Atmospheric forcings are provided with 3- and 24-hourly frequencies by the ERA-interim dataset (Dee et al., 2011) provided by the European Centre for Medium-Range Weather Forecasts.

In order to investigate the effect of sea ice on the topology of hydrodynamic provinces, we also use the product's sea ice concentration fields. The sea ice concentration in each grid cell is defined as the fraction of the cell's area that is covered by sea ice. Sea ice thickness and sea ice velocities are not considered, since the presence of sea ice in our upper layer is implicitly incorporated in our velocity fields.

3.2. Spatial Domain

To study the Arctic, we limit ourselves to an open domain defined above 60°N and only load hydrodynamical data for this area (see Figure S2 in the supporting information). An inherent effect of having a domain with an open boundary is a loss of connectivity information. There is a possibility that at a time scale τ , two geographic regions within our domain may exchange fluid parcels through currents or eddies that (partially) fall outside of the domain. The loss of information is dependent on the time scale τ at which the trajectories of fluid parcels are investigated, and therefore on the distance from the boundary of our domain: the further a parcel is from the boundary, the less probable it is to reach the boundary within our integration time τ , so information is lost less likely.

In particular, our domain choice causes the Denmark Strait and Davis Strait to be disconnected, since the southernmost tip of Greenland lies outside the domain. The East Greenland Current flows around this tip at Cape Farewell where it transitions into the weaker West Greenland Current. This is a clear example of a location where information loss occurs: trajectories between the east and west of Greenland cannot be resolved, so although these areas may be connected by flow, this is not visible in our Lagrangian flow network.

3.3. Domain Discretization

The domain needs to be discretized into bins of comparable shape and size in order for the hydrodynamics to be mapped onto a network. Since we are considering the Arctic, we cannot use a regular grid for this,

since the convergence of the meridians causes bins to have drastically varying areas. To circumvent curvature-related issues, we use a *icosahedral-hexagonal grid*, which is composed of a tessellation of hexagons and 12 pentagons, which are all of similar size and shape. A detailed account of its construction is found in the supporting information (Text S5).

Next to their desirable isotropy, another advantage of icosahedral-hexagonal grids is that adjacent bins always share a border. This is in contrast to regular grids, where the rectangular bins have diagonal neighbors. In our grid, particles leaving a bin always spend some time in directly neighboring bins.

3.4. Particle Simulation

Particle trajectories are calculated through fourth-order Runge-Kutta integration in the *Parcels* Lagrangian framework (version 2.1.2) (Delandmeter & van Sebille, 2019). Even though we consider passive buoyant particles, *Parcels* allows us to specify particle behavior, which is useful for setting the boundary conditions at our open boundary and at the coast. At the open boundary, we freeze particles that reach latitudes below 60°N. Therefore, in our transition matrix, transport to regions outside the domain is represented through the bins that lie at the boundary. We choose this strategy over a leaking system where particles leave the domain, because it is useful to determine the connectivity of the regions of exit with respect to the rest of the domain. Another solution to capture transport outside of the domain is to represent the Atlantic and Pacific basins south of the domain by singular absorbing nodes. However, this would cause a large fraction of all outgoing links to point to these two nodes, which can cause *Infomap* to cluster large geographically separated regions together solely because they share the same absorbing node.

Particles may also get stuck as they get pushed toward cells where their speed becomes zero. This can be the case at land cells, where the meridional and zonal velocity fields become zero. Particles can reach these cells since the velocity fields do not have impermeable boundary conditions at the coast. Although we could specify a boundary condition where particles reaching the coast are sent back into the ocean domain, methods to do so are ambiguous. Rossi et al. (2014) and Ser-Giacomi et al. (2015a) remove these stuck particles. Instead, we keep these stuck particles in $\mathbf{P}(t_0, \tau)$, and interpret this as the beaching of buoyant particles.

We choose a domain discretization using an icosahedral-hexagonal grid at grid level 7. At this refinement level, the average area of a hexagonal bin in the Voronoi diagram is 3,113 km², while the average distance between adjacent bin centers is 60.16 km (0.54°). To assess *Infomap*'s sensitivity to the bin size (and thus to the total number of nodes), we repeat a part of our analysis with icosahedral-hexagonal bins at grid level 6 (see supporting information Text S6). We initialize our particles on the vertices of the triangles of the icosahedral grid at grid level 11, such that distances between particles are about 3.8 km, smaller than the hydrodynamic grid scale. Particles that lie on land are removed, leaving a total of 1,450,665 particles to be simulated. Bins that contain no land initially contain between 253 and 258 particles, the slight variation being due to irregularities in the grid. The number of initial particles in bins that contain land may be much lower. The refinement levels for domain discretization and particle initialization are chosen such as to balance accuracy (a sufficiently fine resolution and large number of particles per bin) with computational efficiency (total number of particles).

To investigate connectivity at different time scales, we simulate particle trajectories for an advection time of 90 days. Particle locations are stored daily, such that connectivity at intermediate time scales can also be assessed. Specifically, we look at $\tau = 30$ and 90 days, which have been used by Ser-Giacomi et al. (2015a), motivated by their ecological relevance.

We choose an advection timestep Δt of 20 min. When comparing the locations of particles released at the same location, but advected with a timestep of 1 min, the average Euclidean distance after 30 days is of the order 3 km. Therefore, we assume that using this advection timestep, we are able to resolve trajectories to a high degree of accuracy. With such a timestep, given horizontal surface velocity magnitudes in the order of 1 ms⁻¹, particles travel approximately 1 km per timestep. Since our bin size is about 60 times larger, this timestep is small enough to prevent particles leapfrogging between different bins.

We release particles at 1 March and 1 September, such that seasonal conditions correspond to high and low sea ice extent respectively. We carry out these simulations for each year between 1993 and 2018, which allows us to find trends and patterns of connectivity that persist over years. For the year 2017, we carry out simulations at the start of each month, in order to investigate seasonal effects.

3.5. Matrix and Graph Construction

The initial and final locations of the simulated particles are used to construct transition matrices $\mathbf{P}(t_0, t_0 + \tau)$ and their corresponding network description, as described in Section 2.1. While the bin sizes and number of initialized particles in our domain discretization vary, these variations are normalized when constructing the transition matrix.

Using the icosahedral-hexagonal grid for a discretization into bins, the domain contains 6,614 bins that (partially) contain fluid, such that $\mathbf{P}(t_0, t_0 + \tau)$ is a square matrix with dimensions 6,614 rows and columns.

3.6. Community Detection Using *Infomap*

Finally, we obtain a division of our network into clusters by using *Infomap* (version 1.0.0-beta.51), which we configure to take into account the characteristics of our flow network.

To start with, we specify that the network should be interpreted as a directed network. In order for the steady-state visiting frequencies to be determined, we use the standard value for the teleportation probability of $\sigma = 0.15$. By making use of the *unrecorded teleportation* scheme, solutions are robust in the regime $\sigma \in (0.05, 0.95)$. For lower values, the steady state visiting frequency π becomes unstable, while for higher values, the steady state approaches the weights of each link (see supporting information Text S2 [Lambiotte & Rosvall, 2012]).

In addition, we make sure that self-edges, which point from a node i to itself, are included. In fact, $\mathbf{P}(t_0, \tau)$ often has values on the diagonal, meaning that particles stay within a bin after timestep τ , making self-edges an indispensable part of our flow description.

Furthermore, we only consider a two-level community description, meaning that we do not consider nested communities. In a nested description, it can be difficult to assess which communities should be expanded, making it hard to compare structures. Instead, we let *Infomap* only return one layer of communities, which are partitioned as to minimize the map equation.

Lastly, different experiments are carried out to determine a Markov-time parameter t_m that produces communities of an appropriate spatial scale. This scale should be large enough to allow for comparison of solutions between different seasons and such that we may attempt retrieve oceanographically relevant structures.

Each time we run *Infomap*, we let it run its outer optimization loop 20 times to obtain different stochastic realizations, after which only the partition \mathbf{P} with minimum $L(\mathbf{P})$ is saved. This ensures that partitions are of a high quality, while keeping computation times reasonably low. Contrary to what Ser-Giacomi et al. (2015a) report for the Mediterranean, solutions do not converge by running *Infomap* more often. Instead, L converges only in the coarse- and fine-tuning steps that *Infomap* executes in one run. The corresponding solution depends on the random order in which nodes are moved by the algorithm, such that different runs do not produce the same result. A higher quality partition may be found by running *Infomap* more often, which can sometimes yield a better value of $L(\mathbf{P})$, but we later show that further improvements in $L(\mathbf{P})$ have only small influences on the global coherence ratio and global mixing parameter.

In all experiments, the same random seeds are used. The seeds are varied when we explicitly aim to compare differences in solutions due to *Infomap*'s stochastic nature (varying the seed value between 1 and 100).

In general, solutions are compared from multiple perspectives. First, the topology is assessed with respect to oceanographic features as well as the persistence of boundaries among different solutions corresponding to different time intervals. Furthermore, we assess whether our criteria for coherence and mixing are met, as

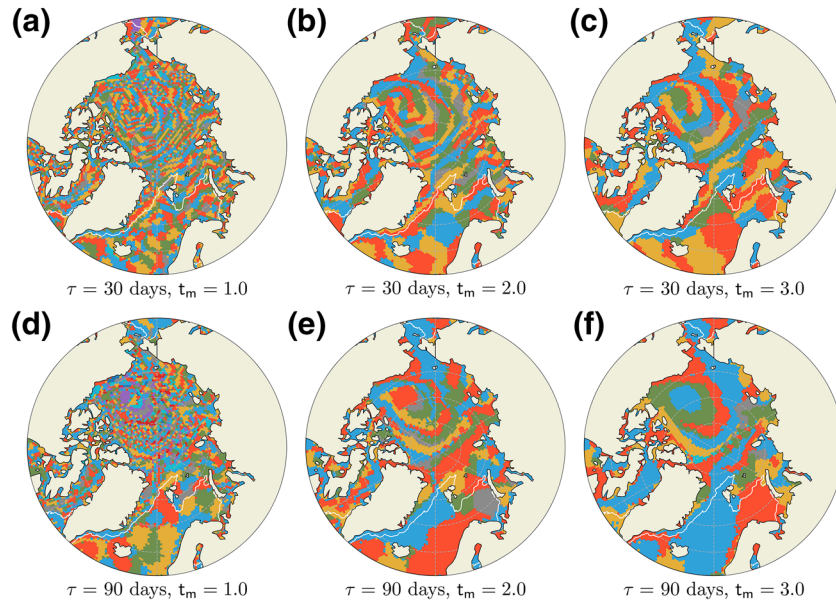


Figure 1. Comparison of solutions returned by *Infomap* for different values t_m for $\tau = 30, 90$ days (a-f). Particles are initialized on $t_0 = 1$ March 2018. White contours indicate the average sea ice extent in March 2018, defined as the contour line corresponding to a sea ice concentration of 15%.

defined in Section 2.3. Lastly, we compare values of $L(P)$ through the map Equation 5. This way, solutions are assessed from an oceanographic and information theoretical perspective, and the connection between these two perspectives is evaluated.

4. Results

Due to the large temporal extent of our dataset and the various configurations possible with *Infomap*, there are many dimensions through which we can study connectivity in the Arctic using the methods we have laid out. For example, to investigate the system at different temporal and spatial scales, we can carry out analyses for a range of values of τ and Markov-times. Also, since 26 years of data are available, in principle we can run *Infomap* over many different transition matrices. However, while aiming to provide a comprehensive overview of the different features that govern the topology of hydrodynamic provinces, we wish to avoid a lengthy and excessive treatment of each variable that could potentially be at play. This motivates carrying out the following experiments: first, we compare solutions for one transition matrix obtained with different Markov-times and we choose one value to carry out all other analyses with. Then, we aim to assess

to what extent solution degeneracy introduces variations in community topology and our quality parameters among different solutions for four transition matrices, corresponding to March and September 2018 and different time scales. Subsequently, we assess the effect of having an open boundary in our domain. Having assessed these effects, we investigate the persistence of community boundaries over time. Lastly, we examine connections between community structures, sea surface velocities and sea ice and investigate temporal trends and seasonal cycles.

Maps with hydrodynamic provinces (Figures 1 and 2) are colored with arbitrary colors in such a way that two neighboring communities never share the same color. However, since *Infomap* does not have any information on how the network is embedded in space, communities can exhibit enclaves, meaning that different parts of the same community may not be connected in space. Due to limitations in visualization, these enclaves are not explicitly indicated in the figures.

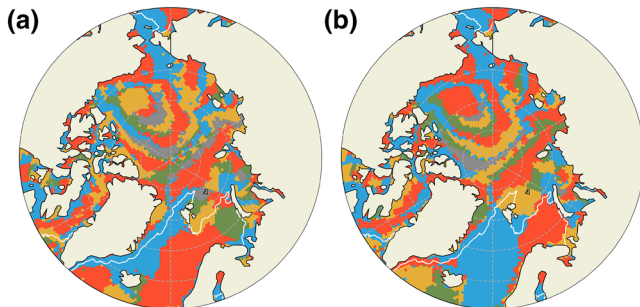


Figure 2. Two solutions found for $P(t_0 = 1$ March 2018, $\tau = 90$ days). White contours indicate average sea ice extent in March 2018. (a) $L(P) = 6.6947$, $\rho_{t_0}^r(P) = 0.7831$, $M_{t_0}^r(P) = 0.3401$ (b) $L(P) = 6.6948$, $\rho_{t_0}^r(P) = 0.7934$, $M_{t_0}^r(P) = 0.3346$.

4.1. Choosing a Markov-Time

We investigate hydrodynamic provinces at different spatial scales by tuning the Markov-time parameter t_m . However, it is a priori not clear what spatial scale should be used. For the sake of consistency, we wish to continue with just one value of t_m . This value should yield solutions with communities at a spatial scale that is, appropriate for analysis for different values of τ . A specific definition of a good spatial scale depends on the specific application. Here, we limit ourselves to two broad criteria. On the one hand, we should be able to assess which bins to flag as boundary bins. In order to distinguish a community's boundary from its interior, the community should at least be three bins wide. Preferably, communities are even wider, since boundary locations can differ among solutions due to degeneracy. If communities are too small, ensembles of solutions become too noisy to analyze visually. On the other hand, communities should not be too large such that they span tens of degrees latitude or longitude, containing many features that are known to function as physical barriers to transport, since this would be at odds with our aim of finding communities with boundaries that correspond to barriers to transport themselves. Both criteria should hold for the range of time intervals $\tau = 30\text{--}90$ days.

Figure 1 shows a comparison of solutions returned by *Infomap* for different values of t_m and τ . There are a few similarities between the different solutions. For example, all solutions exhibit a circular configuration of hydrodynamic provinces around the Beaufort Gyre. For $\tau = 30$ days, solutions contain filamental structures around the East Greenland Current, indicating that it acts as a barrier in the cross-current direction.

Community sizes increase with an increase in τ or t_m . This is to be expected, since an increase in τ allows Lagrangian particles to travel farther, thus connecting more bins. Moreover, due to the chaotic nature of ocean flow, particles originating from the same bin may over time follow different currents and eddies, such that when longer time spans are considered, the spread in particle distributions becomes larger. From a network perspective, this decreases the average distance between nodes and a random walker on the network can therefore traverse larger distances, connecting bins that are separated by larger distances too. For increasing t_m , random walkers may also traverse more nodes before having their position recorded. Therefore, it makes sense that *Infomap* draws community boundaries at larger distances as either τ or t_m increases.

For $\tau = 30$ days and $t_m = 1.0$ (*Infomap*'s default value), the solution consists of many small communities that have sizes that are too small to assess boundary persistence. For this t_m and both values of τ , communities are especially small in the presence of sea ice. This makes sense, as the surface velocities in these areas are drastically lower. In these areas, some communities also consist of many spatially disconnected bins. Communities are largest for $\tau = 90$ days and $t_m = 3.0$. For this solution, one hydrodynamic province spans from the edge of the domain at 60°N between Iceland and Norway all the way to the sea ice boundary and protrudes far into the sea ice. We deem the value $t_m = 3.0$ too high, since surface velocities at the sea ice boundary drop drastically and it should therefore provide a natural boundary in the system. We choose to continue with $t_m = 2.0$, since the solutions for both values of τ exhibit a community scale that fits both of our broad criteria.

4.2. Individual Solutions and Solution Degeneracy

As discussed in Section 2.2, *Infomap* suffers from solution degeneracy due to its stochastic and heuristic nature. In order to use *Infomap* to study oceanographic structures, the persistence of boundaries, and temporal trends, we must first evaluate the role that solution degeneracy plays. We do this by running *Infomap* on the same transition matrix with exactly the same parameters, only varying the random seed. We compare results for 100 different seeds in terms of codelength, global coherence, global mixing and boundary persistence. Differences in results are then only due to degeneracy.

For 100 solutions obtained for $\mathbf{P}(t_0 = 1 \text{ March } 2018, \tau = 90 \text{ days})$, the average codelength is $L(\mathbf{P}) = 6.694$, while the associated standard deviation is 0.011 (from now on reported in parentheses). The average global coherence ratio is $\rho_{t_0}^r(\mathbf{P}) = 0.7843 (\pm 0.0067)$, while the average global mixing parameter is $M_{t_0}^r(\mathbf{P}) = 0.3392 (\pm 0.0023)$.

Figure 2 shows the two solutions that have their codelengths closest to the average value. Both solutions are good solutions since they have been obtained by running *Infomap* for each different seed 20 times and

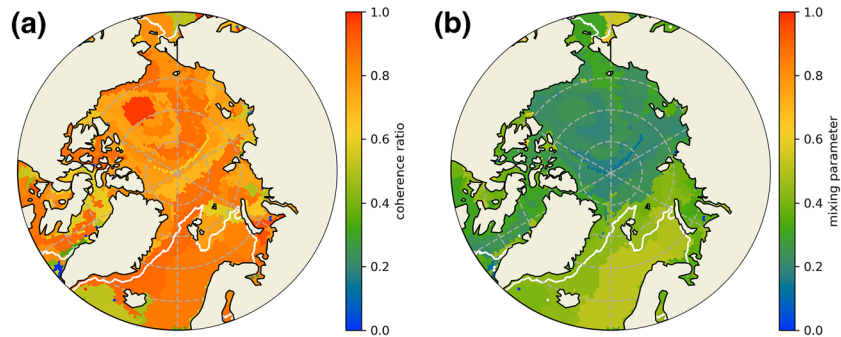


Figure 3. The coherence ratio (a) and mixing parameter (b) associated to each community in the partition depicted in Figure 2a, for $P(t_0 = 1 \text{ March } 2018, \tau = 90 \text{ days})$.

picking the partition with the lowest codelength. The codelengths of the two solutions differ by 0.0001. Both solutions exhibit similar topologies in the Davis Strait, the Beaufort Gyre, and the Chukchi Sea. However, for certain areas, topologies are very different. This can clearly be seen in the Norwegian Sea. The solution in Figure 2a separates the Norwegian Sea from the Greenland Sea, while the solution in Figure 2b clusters these seas together.

Coherence ratios and mixing parameters exhibit spatial patterns. This can be seen in Figure 3, which shows the coherence ratio and mixing parameter associated to each community of the partition depicted in Figure 2a. Communities that lie close to the boundary of the domain can exhibit low coherence ratio since particles may exit the domain from these communities. Bins where particles exit the domain are often clustered as single communities. In contrast, the community at the center of the Beaufort Gyre shows a coherence ratio close to 1, meaning this community retains almost all its particles. This is because around the center of the Beaufort Gyre, the flow is anticyclonic with little transport in the radial direction, preventing particles from leaving the center.

Mixing parameters are generally higher for communities in ice-free regions, likely due to higher velocities and the presence of eddies that can stir Lagrangian particles across a community. Values are especially high in the Norwegian Sea, which contains the Norwegian Current which exhibits baroclinic instability (Mysak & Schott, 1977). This area is characterized by high mesoscale activity (Hansen et al., 2010). In regions with sea-ice, transport is mostly in the anticyclonic direction, preventing particles from efficiently mixing within their communities. The solution depicted in Figure 2b exhibits similar patterns in coherence and mixing (see Figure S6 in the supporting information).

Ser-Giacomi et al. (2015a) investigate the persistence of community boundaries across years. We follow a similar procedure to assess the persistence at which *Infomap* draws boundary locations. We flag the bins that lie at the interface between two communities as a boundary bin and assess the frequency at which each bin is marked as such among our 100 solutions. This is shown in Figure 4 for March and September 2018, with $\tau = 30$ and 90 days. We repeated this analysis with icosahedral-hexagonal bins at grid level 6 (see supporting information Text S6) and found similar persistent boundaries. However, grid level 7 reveals structures at a finer scale.

Infomap draws the most boundaries in regions with sea ice. A circular structure around the anticyclonic Beaufort Gyre can clearly be seen, coinciding with the underlying anticyclonic flow and the flow of sea ice. The circular structure of persistent boundaries reaffirms the notion that flow in the radial direction is mostly prohibited. Boundaries are also persistently found separating the Irminger Basin and Iceland Basin from each other and the rest of the domain. These basins are physically separated by the Irminger Current, coinciding with the Reykjanes Ridge. A boundary also persists at the edge of the continental shelf east of Greenland, where the East Greenland Current is located. For $\tau = 30$ days, small communities persist in the Norwegian Sea, but locations differ between March and September. Especially for $t_0 = 1 \text{ September}$, boundary-free regions can clearly be seen. For $\tau = 30$ days, the Norwegian Sea contains ring-like boundaries, with radii of the order of 200 km. Their boundary occurrence is between 0.7 and 1. For $\tau = 90$ days *Infomap* is

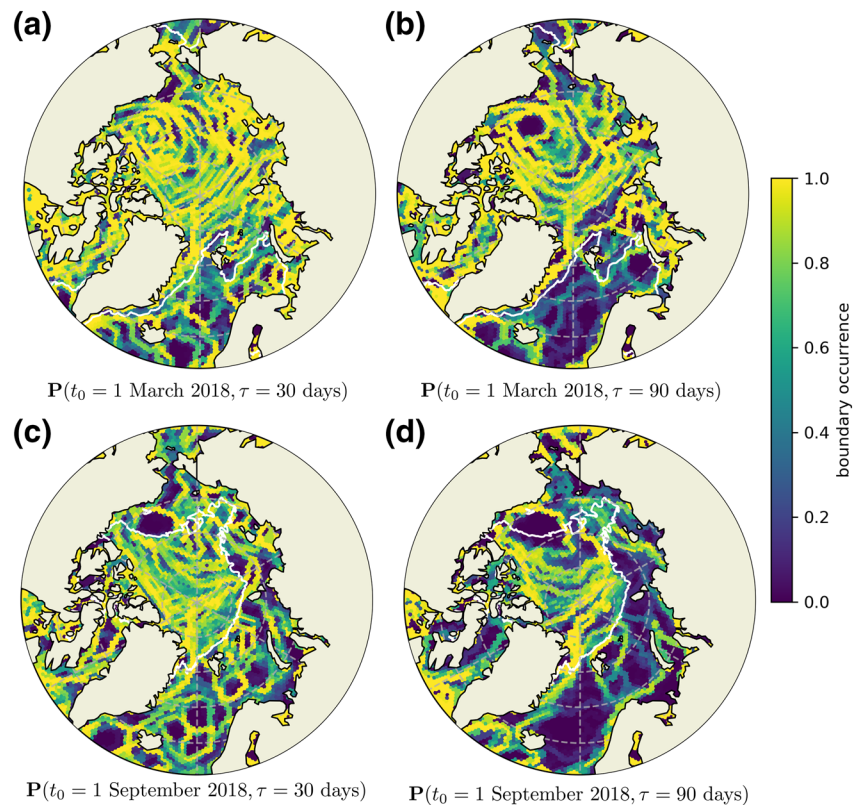


Figure 4. Persistence of community boundaries in a set of 100 solutions found for different transition matrices (a-d). White contours indicate average sea ice extent in March and September 2018, respectively.

less persistent in drawing boundaries in and between the Norwegian Sea and Greenland Sea. For this τ , a frontier is visible running from the coast of the Scandinavian peninsula to Novaya Zemlya, isolating the White Sea and its outflow.

Next to plotting where *Infomap* finds persistent boundaries, it is also insightful to investigate which cells are persistently clustered together by *Infomap*, as this is not revealed by the persistence of boundaries (Calatayud et al., 2019). This is further explored in the supporting information (Text S7).

4.3. Sensitivity to Domain Boundary

Having a domain with an open boundary may entail losing information about connectivity between bins, since trajectories may also include locations outside the domain. Since particle trajectories are frozen as soon as a particle exits the domain, the choice of the domain boundary is of influence on the transition probability of a particle. We assess the extent to which our community boundaries are influenced by the latitude at which we define our domain. The codelength of a partition yielded by *Infomap* depends on the transition probabilities between bin pairs. Changes in transition probabilities due to shifting the domain boundary should be localized to bins that lie close to the boundary. Therefore, we expect the influence of our domain choice on the locations of boundaries to be localized at the domain boundary.

We investigate the influence of the location of the domain boundary by comparing the boundary persistence for different domains, but with simulations obtained for the same values of t_0 and τ . We choose $t_0 = 1$ September 2018 and $\tau = 90$ days, such that communities are large enough for the boundary persistence not to be too noisy and to reduce the effect of sea ice. Ideally, we would compare to the case where we do not have an open domain at all, which can be achieved by expanding the domain to the global ocean. However, this would come with increased computational costs, partly due to loading significantly more hydrody-

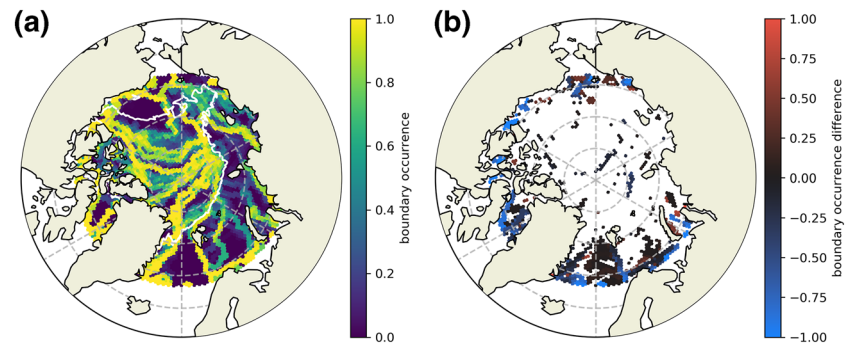


Figure 5. (a) Persistence of community boundaries for 100 solutions obtained for $\mathbf{P}(t_0 = 1 \text{ September } 2018, \tau = 90 \text{ days})$, with the redefined domain being bounded by 70°N . White contours indicate average sea ice extent in September 2018. (b) Difference in community boundary persistence, obtained by subtracting the persistence in the 70°N domain in (a) from the persistence in the 60°N domain in Figure 4d. Differences are only shown when they exceed the standard deviation associated to flagging a bin as a boundary across the 100 solutions obtained using the 70°N domain (calculated for each bin).

namical data and also due to the extra computation of trajectories of particles below 60°N . Instead, we compare our normal domain bounded at 60°N , for which the boundary persistence is found in Figures 4d, to a smaller domain bounded at 70°N . Again, we freeze particles that reach the boundary at 70°N . Since particle trajectories are obtained deterministically by using Equation 1, trajectories of particles that do not reach latitudes lower than 70°N are the same as when considering a domain bounded by 60°N . Large parts of the resulting transition matrix should thus be equal to that of 60°N . We apply *Infomap* 100 times on the transition matrix obtained for the simulation in the modified domain, which for 67% of originating bins (columns in the transition matrix) is completely equal to the transition matrix in the original domain. The persistence of boundaries in these 100 new solutions is shown in Figure 5a.

Infomap may either draw community boundaries in the two domains at different locations due to the underlying transition matrices being different, either due to particle trajectories not being included or being cut short, or it may do so due to its heuristic and stochastic nature. Figure 5b shows the difference in persistence of solutions obtained from the 60° and 70°N domains. Here, bins are left white if this difference is lower than the standard deviation associated to flagging a bin as a community boundary across the 100 solutions obtained using the 70°N domain. This allows us to see where differences in community boundaries are likely due to solution degeneracy and where they may be due to the different domain choice. Note that the differences in boundaries are significant mostly near the domain boundary. A notable exception is in the Greenland Sea, where this difference comprises a large portion of bins. However, differences in the persistence of community boundaries here are still small. In general, persistent community boundaries in the interior are not located differently. Major differences are localized to the domain boundary. We theorize that when comparing communities in the 60°N domain to communities obtained from transition matrices for the global ocean, differences in boundary persistence should similarly be localized to our current domain boundary at 60°N . We thus assume that boundary persistence values as found in Figure 4 are mostly the same as when instead we would have considered the global ocean, with differences mainly being localized to the vicinity of the domain boundary.

4.4. Persistence of Transport Barriers Over Time

It is also insightful to assess the persistence of boundaries between hydrodynamic provinces over different years. Like Rossi et al. (2014) and Ser-Giacomi et al. (2015a), we take the average of the community boundaries of solutions obtained from transition matrices corresponding to different years and seasons. However, we take solution degeneracy into account by including 100 solutions for each transition matrix.

Figure 6 shows the persistence of boundaries averaged over the years 2009–2018 for March and September, with $\tau = 30$ and 90 days. Each subfigure is thus composed using 10 transition matrices, corresponding to

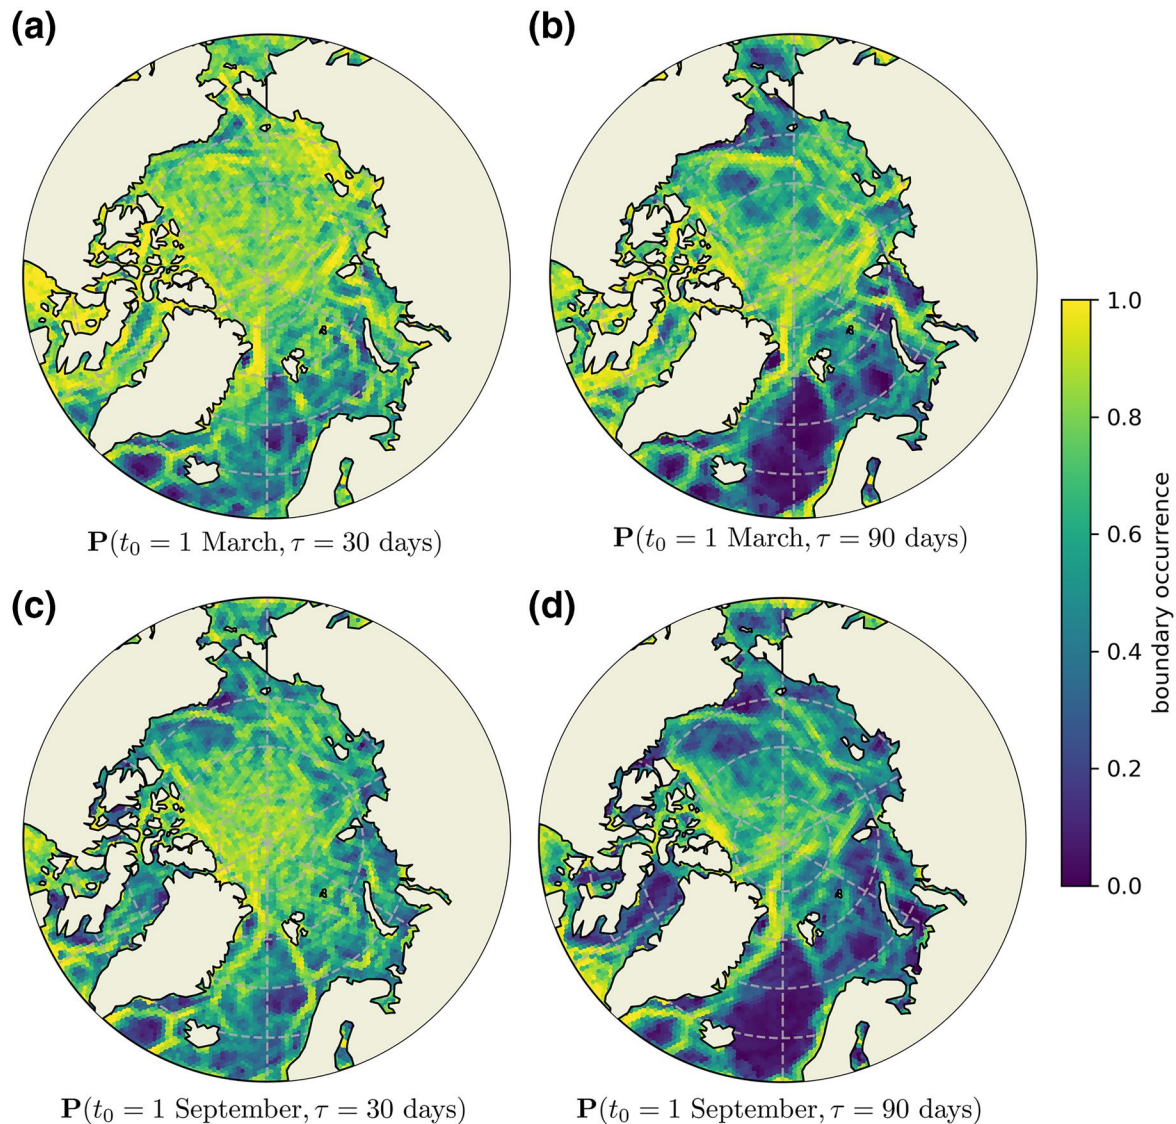


Figure 6. Persistence of community boundaries between 2009 and 2018. For each transition matrix (a-d), 100 different solutions are obtained using *Infomap*.

each year, and for each transition matrix, 100 solutions are obtained using *Infomap*. This way, boundaries that are due to degeneracy or natural variability are filtered out.

Across all solutions, we again observe a circular structure around the Beaufort Gyre. For particles released in September, the East Greenland Current is also persistently visible. The North Atlantic Current, including the Norwegian Current and West Spitsbergen Current are prominently visible for particles initialized in September with $\tau = 30$ days, and to a lesser extent for the solutions with $\tau = 90$ days. The Irminger Current persists across solutions. For $\tau = 90$ days, the solutions for particles initialized in March show a boundary at the north-eastern coast of Iceland, while such a boundary is not visible for particles initialized in September. This may be due to the seasonality in the strength of the North Icelandic Irminger Current (Logemann & Harms, 2006). For $\tau = 30$ days, boundaries occur more often in the Norwegian Sea and Greenland Sea than for $\tau = 90$ days, with boundaries not being located persistently. We think these boundaries are related to the high mesoscale activity in this region, instead of being due to persistent currents. At shorter time scales, *Infomap* finds smaller clusters in these seas, while for $\tau = 90$ days, particles are sufficiently stirred for these clusters to coalesce.

4.5. Correlations to Sea Surface Velocities and Sea Ice Concentrations

To better understand how *Infomap* is affected by the physics that give rise to our transition matrices, we investigate and attempt to explain correlations between codelengths, community boundaries, coherence, mixing, surface velocities and sea ice concentrations. Of these quantities, codelengths and community boundaries pertain to the graph description of our physical system, while surface velocities and sea ice concentrations are inherent to the physical fields that govern the dynamics of Lagrangian particles. Coherence and mixing are dependent on both the community division found by *Infomap*, as well as the physical trajectories of Lagrangian particles, thus bridging the physical and graph descriptions of our system. We note that sea ice concentrations are not explicitly used to determine particle trajectories, but the presence of sea ice does influence the surface velocity field and therefore implicitly affects the system. In our full time series (1993–2018), sea ice concentration and the seawater velocity magnitude are anticorrelated, with a Pearson correlation coefficient of $r = -0.36$ and an associated p -value of 0 (below machine precision).

To investigate the correlation between codelengths, global coherence ratios and global mixing parameters, we use a set of 1,000 solutions, comprised of 100 solutions for 10 transition matrices obtained through simulations between 2009 and 2018 with $t_0 = 1$ September and $\tau = 90$ days. For these global quantities, we find that the codelength and global mixing parameter have a negative Pearson correlation coefficient of $r = -0.093$, with an associated p -value of $p = 3.10 \times 10^{-3}$. Simultaneously, we find a positive correlation between codelength and global coherence, with $r = 0.18$ and an associated p -value of $p = 1.40 \times 10^{-8}$. While these correlations have a high statistical significance, they are also weak, and no clear relation can be established between codelength, global coherence and mixing.

To locally assess correlations between the velocity, sea ice, coherence, mixing, and boundary persistence of each bin, we make use of the same 1,000 solutions for the transition matrices $\mathbf{P}(t_0 = 1$ September, $\tau = 90$ days) between 2009 and 2018. We use transition matrices with particles released in September, since a larger portion of the domain is ice-free. The corresponding boundary persistence can be found in Figure 6c, while the other quantities can be found in Figure 7.

Meridional and zonal velocities are averaged per grid cell over September, October and November 2018, interpolated onto the icosahedral-hexagonal grid, and converted into mean speed. Sea ice concentrations are similarly averaged and interpolated.

We find a positive correlation between mean speed and boundary persistence, with a Pearson correlation coefficient of $r = 0.38$ and $p = 0$ (below machine precision) in regions where the sea ice concentration is less than 0.15. If we include all bins, this statistically significant correlation vanishes. This indicates that in ice-free regions, currents indeed correlate to community boundaries and thus provide barriers to transport. This correlation disappears in the presence of sea ice, meaning that in the sea ice regime, other factors govern the existence of boundaries.

When comparing correlations with sea ice, we find a weak positive correlation between coherence ratio and sea ice concentration of $r = 0.21$ and $p = 6.30 \times 10^{-69}$. On visual inspection of the average coherence ratio in Figure 7c, this relation to sea ice is difficult to discern. The correlation between sea ice and coherence ratio may be biased due to the low coherence ratio of some communities at the edges of the domain, where particles may escape to communities containing only a few bins.

The mixing parameter and sea ice concentration are negatively correlated with $r = -0.24$ and $p = 2.60 \times 10^{-88}$. From Figure 7d we can observe that mixing is generally higher in ice-free regions. However, the mixing parameter is also low around the East Greenland Current. This makes sense, since this current flows southward, such that within communities in this region, it is only possible for particles to spread to bins that lie south. This is similar to the mainly anticyclonic transport around the Beaufort Gyre. In contrast, mixing is strong in the Norwegian Sea and Barents Sea. The high mesoscale activity in the Norwegian Sea may provide relatively efficient mixing in the communities located there.

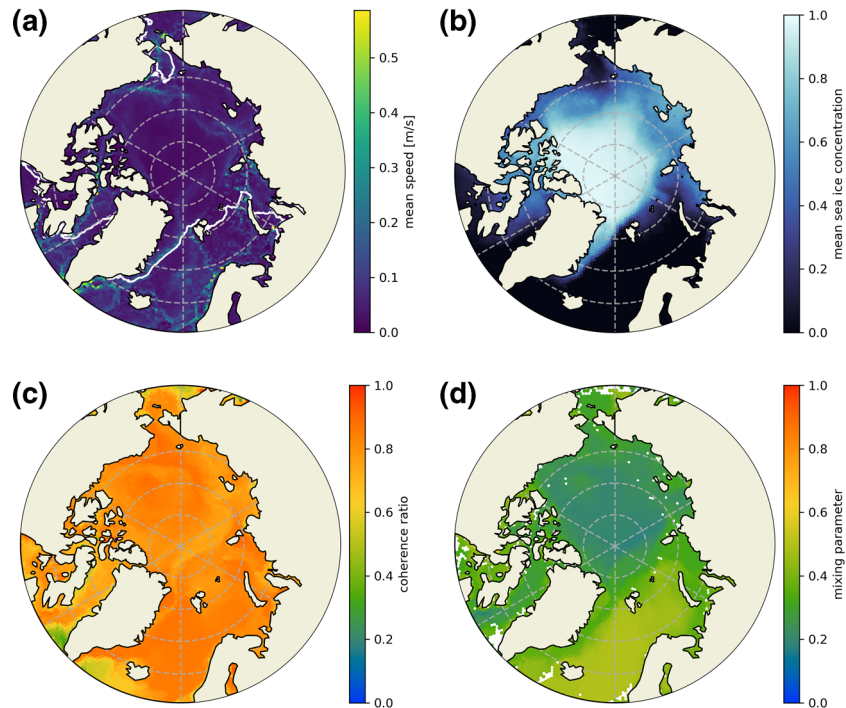


Figure 7. Average speed, sea ice concentration, coherence ratio and mixing parameter for each bin. Coherence ratio and mixing parameter correspond to average values of the communities a bin is partitioned with in each of the 1,000 solutions for $P(t_0 = 1 \text{ September}, \tau = 90 \text{ days})$ between 2009 and 2018. (a) Mean speed in September-October-November 2009–2018, interpolated on our icosahedral-hexagonal grid. (b) Mean sea ice concentration in September-October-November 2009–2018, interpolated on our icosahedral-hexagonal grid. (c) Average coherence ratio per bin. (d) Average mixing parameter per bin.

4.6. Trends and Seasonality

Since the sea ice extent in the Arctic experiences a strong seasonal variation, which has implications for ocean surface flow, we expect the quality of solutions to be affected by this. Furthermore, the decrease in summer sea ice extent that has been observed in the past decades is also of influence on surface flow, thus it also affects solutions found by *Infomap*. We assess these effects by looking at the yearly and monthly temporal evolution of solution quality.

4.6.1. Annual Trends

Figure 8 shows the evolution of sea ice, codelength, global coherence ratio and global mixing parameter calculated from 100 degenerate solutions obtained for $\tau = 90 \text{ days}$ and $t_0 = 1 \text{ September}$ for each year between 1993 and 2018. Estimated trends are included, based on linear regression. For codelengths, the standard deviation associated to the solution degeneracy is much smaller than the differences in average codelengths between different years, and is thus hardly visible. This indicates that most differences among solutions cannot be attributed to solution degeneracy, but are instead due to differences in the underlying flow, mirrored in the transition matrices. The global mixing parameter also exhibits standard deviations that are smaller than the variation of mean values between years. In contrast, for the global coherence ratio, the standard deviations due to solution degeneracy are of the order of the variation of mean values between years.

We find a negative correlation between sea ice area and codelength, with $r = -0.44$ and $p = 0.023$. The correlation between sea ice and the global mixing parameter is also negative, with $r = -0.75$ and $p = 9.40 \times 10^{-6}$. We do not find a significant correlation between sea ice and the global coherence ratio.

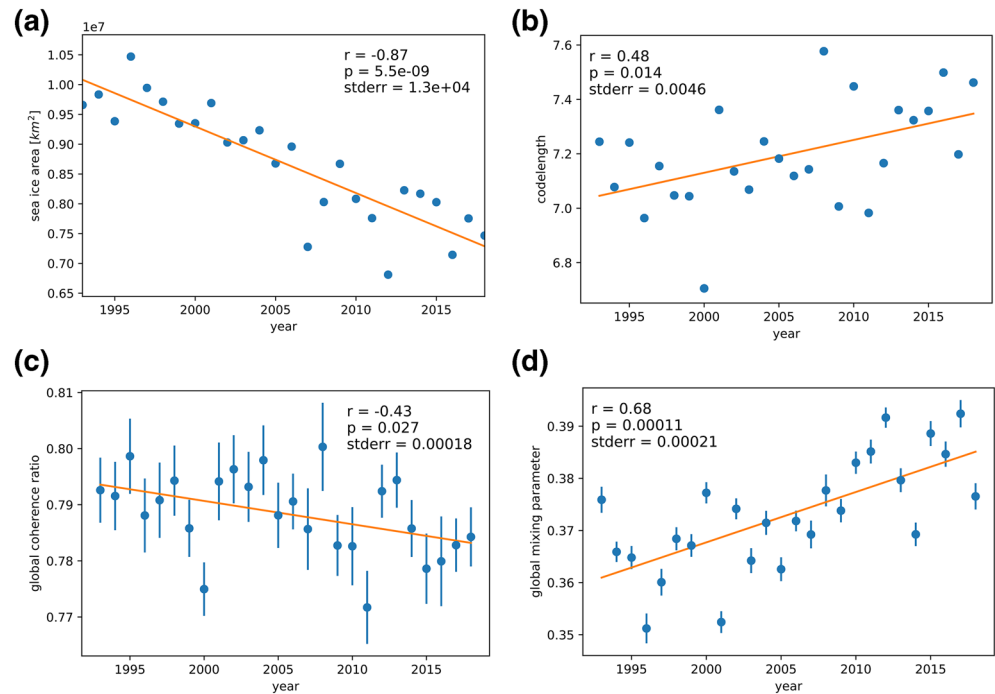


Figure 8. Evolution of sea ice, codelength, global coherence ratio and global mixing parameter in September between 1993 and 2018. One hundred solutions have been obtained for each year by initializing particles on the first day of 1 September. $\tau = 90$ days. Trend regression is indicated in orange, including associated correlation coefficient r , p -value (corresponding to a hypothesis test with the null hypothesis corresponding to a zero slope, using the Wald Test), and standard error. Standard deviations related to solution degeneracy are indicated using error bars, except for sea ice area. (a) Mean sea ice area. (b) Codelength. Standard deviations are on average 0.012, making error bars invisible. (c) Global coherence ratio. (d) Global mixing parameter.

While the sea ice area exhibits a clear downward trend, the codelength and global mixing parameter show positive trends, increasing as the Arctic sea ice cover shrinks. The coherence ratio shows a slight downward trend, although with a higher p -value than the codelength, sea ice area and mixing.

To supplement the inspection of solution quality in different years, Figure 9 shows the difference of the coherence ratio and mixing parameter for the period 2009–2018 (as in Figures 7c and 7d) and 1993–2002 (Figure S7 in the supporting information). These time spans correspond to the first and last 10 years of our dataset. No clear spatial pattern can be seen for changes in the average coherence ratios. An increase in the mixing parameter occurs primarily in the East Siberian Sea, Beaufort Sea, Laptev Sea, Kara Sea and Chukchi Sea. These seas are where the summer sea ice loss trends are highest (Meredith et al., 2019). As sea ice is anti-correlated with surface velocity magnitude, we theorize that the loss in sea ice allows the stronger surface flow velocities to provide more efficient stirring of Lagrangian particles.

4.6.2. Seasonal Effects

Seasonal development of the sea ice area, codelength, global coherence ratio and global mixing parameter are assessed by comparing 100 solutions for 12 transition matrices, for which t_0 equals the first day of each month in 2017, while $\tau = 90$ days. Figure 10 shows the monthly evolution of these parameters. For the codelength and mixing parameter, a clear seasonal cycle can be observed, with maxima in summer and minima in winter, which coincides with the seasonal cycle in sea ice area. Indeed, we find a negative correlation between sea ice area and codelength of $r = -0.85$ with $p = 4.00 \times 10^{-4}$ and a negative correlation between sea ice area and the global mixing parameter of $r = -0.85$ and $p = 5.40 \times 10^{-4}$. A seasonal cycle for the global coherence ratio cannot be inferred from the data and we do not find a correlation between sea ice and coherence ratio here. Changes in the average global coherence ratio are small compared to the standard deviation due to the solution degeneracy. A possible explanation for this is that *Infomap* indirectly optimizes for coherence of particles

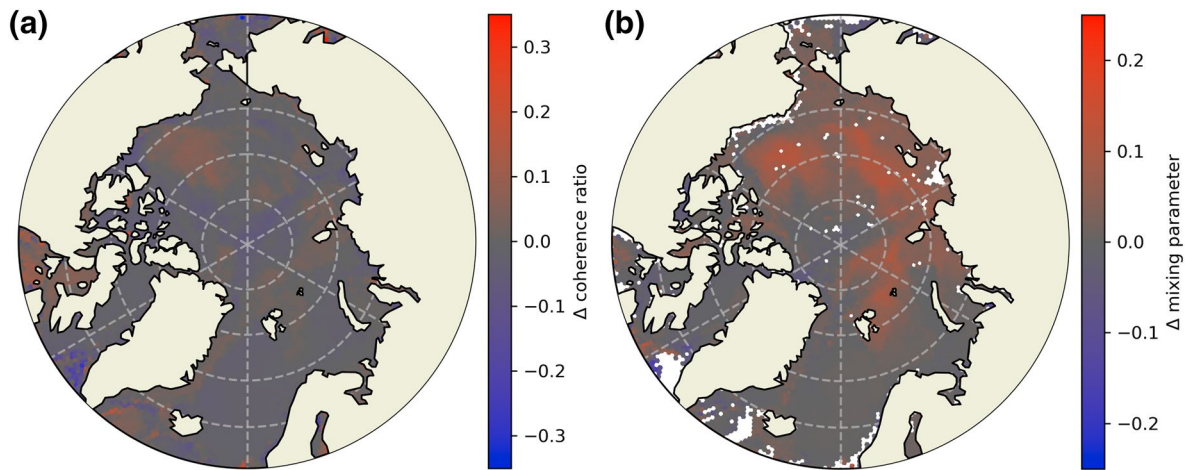


Figure 9. Difference of average coherence ratio and mixing parameter for $\mathbf{P}(t_0 = 1 \text{ September}, \tau = 90 \text{ days})$ between 2009 and 2018 (Figures 7c and 7d) and 1993–2002 (Supporting information Figure S7). (a) Difference in coherence ratio. (b) Difference in mixing parameter.

within its communities. Although the flow changes with the seasons, *Infomap* still finds communities that manage to retain their particles to a similar degree.

Since sea ice coverage is minimal in summer, the seasonal cycle of the codelength agrees with the yearly trend of increasing codelength as sea ice declines. This also corresponds to what we observe for the mixing parameter, which is lower in areas with sea ice and which globally increases over time, as sea ice cover decreases. We interpret this as follows: as the sea ice extent declines, surface velocities increase, allowing particles to travel larger distances, which increases the connectivity between bins. This in turn reduces the distance between edges in the network. This allows a random walker to traverse the network more easily. Nodes that were previously visited infrequently, increase in steady-state visiting frequency, making the distribution of π more balanced. This causes the average codelength to increase, as infrequently visited nodes that are assigned larger codewords are visited relatively more frequently. Simultaneously, in areas covered by ice, mixing is lower than average. As sea ice disappears and velocities increase, these regions become more mixed.

5. Discussion and Conclusions

We have successfully applied the *Infomap* algorithm to detect hydrodynamic provinces in the Arctic Ocean surface by using a network description of the flow. In the context of hydrodynamic provinces, each community partition returned by *Infomap* corresponds to a local minimum of the map equation. The standard deviation of the codelength due to solution degeneracy is small compared to the differences in codelength due to seasonal and yearly variations in flow. Therefore, each individual solution is good in the sense that it corresponds to a low average codelength and the corresponding community partition should have boundaries to transport such that the transitions between different communities are locally minimized. When investigating individual degenerate solutions, such as in Figure 2, different communities in the sea ice free domain can be seen to correspond to different seas. Boundaries have been shown to correlate with velocities, meaning that currents provide effective barriers to cross-community particle exchange. This relation is not limited to a statistical correlation: major currents have been shown to coincide with community boundaries. However, different partitions may each resolve different physically relevant boundaries. To obtain a more complete picture, it is thus useful to consider an ensemble of solutions, such as in Figure 4. These solutions only arise from a single transition matrix, whereas seasonal pictures can be obtained by considering ensembles over different years (Ser-Giacomi et al., 2015).

The boundaries in regions with sea ice seem to arise from the flow slowly moving in an anticyclonic fashion around the Beaufort Gyre, coinciding with the flow of sea ice. Particles move in concert along this mostly laminar, concentric flow and do not travel much in the radial direction. This causes a low mixing in this region. These factors make the community division exhibit a circular or spiral-like structure. Community boundaries

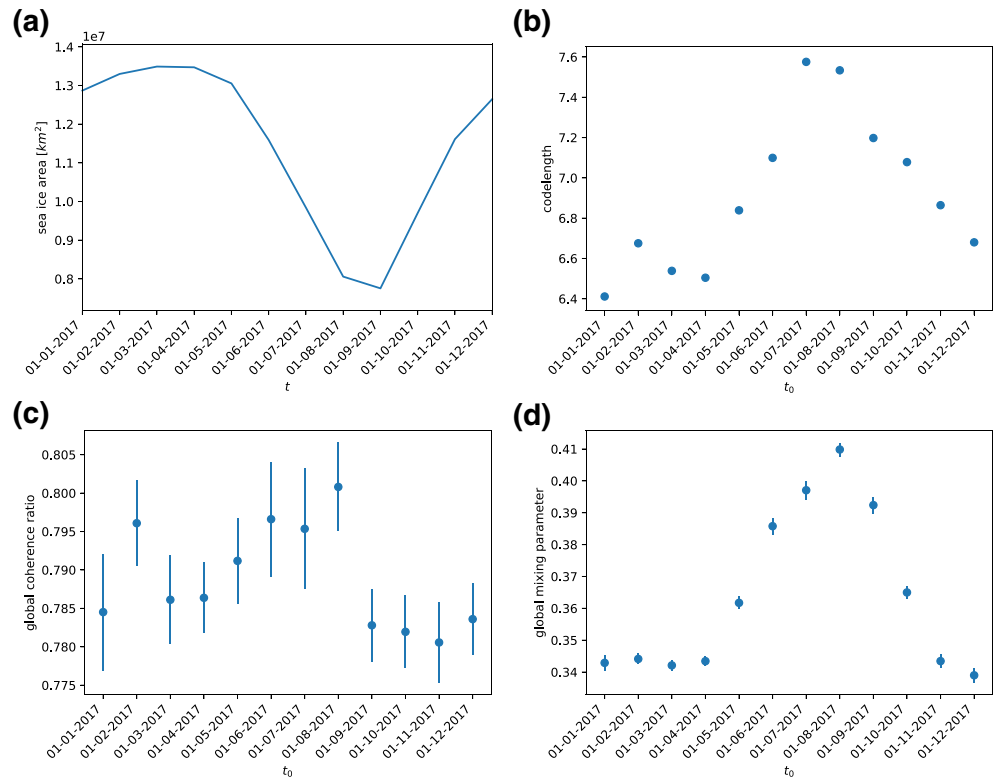


Figure 10. Monthly evolution of the sea ice area, codelength, coherence ratio and mixing in 2017. One hundred solutions have been obtained for each month by initializing particles on the first day of each month. $\tau = 90$ days. Error bars indicate the standard deviation due to solution degeneracy. (a) Mean sea ice area. (b) Codelength. Standard deviations are on average 0.010, making error bars invisible. (c) Global coherence ratio. (d) Global mixing parameter.

experience a higher degeneracy in sea ice free regions. Sea ice cover is shown to be anti-correlated with codelength and global mixing, both seasonally and yearly.

For individual partitions, we note that it is always important to consider the mixing parameter and coherence ratio of the corresponding communities. By only looking at the topology of a community, it may be tempting to assume that any two bins that fall under the same community exchange particles with one another. In contrast, in certain communities, the underlying flow may have one clear direction, such that the corresponding nodes in the network are not strongly connected. This is the case for example, in communities that coincide with strong currents, such as the East Greenland Current. In these communities, particles generally only travel southward, following the flow. It is also the case for the Beaufort Gyre, where flow is restricted to the anti-cyclonic direction.

In many respects, this article extends on the method for identifying hydrodynamic provinces through community detection in Lagrangian flow networks as proposed by Rossi et al. (2014) and Ser-Giacomi et al. (2015a). Foremost, we assessed the role that solution degeneracy plays in yielding different partitions and stress the importance of considering ensemble solutions. Other key differences are the consideration of the Markov-time parameter, the application to a larger, open domain, an assessment of the evolution of global quality metrics in a seasonal and yearly context, and the establishment of correlations between hydrodynamic province boundaries, flow speed and sea ice.

Studying flow by using community detection on Lagrangian flow networks may warrant caveats related to the limited representation of flow as a network, or related to the community detection algorithm. First, the representation of true flow in terms of Lagrangian flow networks is limited by resolution. Especially in the Arctic, mesoscale structures are often not yet fully resolved due to the small Rossby radius in many Arctic regions. Even when such features are resolved, the network representation of the flow only captures these structures statistically. Additionally, since there is a limited amount of trajectories originating from each

bin, representation can be improved by increasing the number of Lagrangian particles that is, simulated. However, this bears extra computational costs, especially when the grid resolution is increased. Furthermore, the representation of the flow is influenced by having an open boundary. When choosing an open domain, it is important to assess to which extent this influences community topologies. For our domain, this effect seems to be localized to the domain boundary, and *Infomap* is still able to find barriers to transport that carry physical significance. Nevertheless, this limitation can only be fully overcome by considering flow in the global ocean.

Our results are also sensitive to multiple parameters. Naturally, the communities returned by *Infomap* are dependent on t_0 and τ , since these parameters govern the time and time scale at which the flow is recorded into a transition matrix. In addition, results are sensitive to the choice of the Markov-time parameter. Here we choose one specific Markov-time parameter to tune *Infomap* in such a way that it provides an appropriate spatial scale for our analyses. For more specific applications, it may be difficult to assess which Markov-time should be considered, since it is impossible to know a priori which Markov-time corresponds to the spatial scale at which one wishes to study connectivity.

Here, we considered the community detection algorithm *Infomap* due to successful previous applications and since its underlying algorithm optimizes a balance between high internal connectivity and good coherence, while emphasizing the flow description of a network. However, the way in which *Infomap* balances coherence and mixing cannot be set explicitly.

Furthermore, the degeneracy of solutions makes the interpretation of a single solution misleading. If the optimal solution with the minimum average codelength could be found, it should in principle yield a partition which optimizes our criteria for strong internal mixing and good community coherence. However, this solution would have been obtained through cumulative approximations of the flow, for example, arising from uncertainties in the observed flow fields, limited spatial and temporal resolutions, a limited amount of modeled Lagrangian trajectories, and freedom in parameter choices. Therefore, there is no reason to assume that the optimal solution corresponds to a division in communities that carries the most physical meaning. This gives further motivation to always consider an ensemble of solutions. Even when an ensemble of solutions is considered, community detection can be useful, but should be interpreted with caution due to the propagation of uncertainties, approximations, and errors.

Because the Arctic domain is subject to climate change, we suggest further research into future connectivity in the Arctic Ocean by using velocity field output from coupled global climate models. This will reveal how changes in the climate affect barriers to transport. Another suggestion is to investigate connectivity in three-dimensional flows. Regions of up- and down-welling respectively, give rise to divergence and convergence in the two-dimensional surface flow field. Convergence of the flow field is mirrored in the Lagrangian flow network by nodes having more incoming edges. In contrast, the three-dimensional velocity field is (by approximation) divergence free, which has implications for the distribution of incoming edges in Lagrangian flow networks, making it more uniform. It is also interesting to investigate barriers to three-dimensional transport in the Arctic, especially in a changing climate, given its relevance to the Atlantic Meridional Overturning Circulation. Lastly, as connectivity has ecological relevance, ecological and marine spatial planning research may benefit from the present study and our methods may be further refined for applications in those fields.

Acknowledgments

Daan Reijnders was supported through funding from the Netherlands Organization for Scientific Research (NWO), Earth and Life Sciences, through project OCENW.KLEIN.085. Erik van Sebille was supported through funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation program (grant agreement 715386). We thank David Wichmann for providing insightful discussions and feedback. We are grateful to Reik Donner, Erick Fredj and an anonymous reviewer for their valuable comments and suggestions. We also thank Martin Rosvall for providing support with *Infomap* and Louis Moresi for his help with constructing the icosahedral-hexagonal grid.

Data Availability Statement

Transition matrices, corresponding networks and community data are available under a Creative Commons Attribution 4.0 International license and annotated code for this work is available under the MIT license (both are available through <https://doi.org/10.24416/UU01-IN00U9>).

References

- Andrello, M., Mouillot, D., Beuvier, J., Albouy, C., Thuiller, W., & Manel, S. (2013). Low connectivity between Mediterranean marine protected areas: A biophysical modeling approach for the dusky grouper *Epinephelus marginatus*. *PLoS One*, 8(7), e68564. <https://doi.org/10.1371/journal.pone.0068564>
- Brin & Page (1998) should be moved to the references of the supporting information and set in alphabetical order after "Bohlin et al (2014)."

- Broekhuizen, N. (1999). Simulating motile algae using a mixed Eulerian-Lagrangian approach. *Journal of Plankton Research*, 21(7), 1191–1216. <https://doi.org/10.1093/plankt/21.7.1191>
- Calatayud, J., Bernardo-Madrid, R., Neuman, M., Rojas, A., & Rosvall, M. (2019). Exploring the solution landscape enables more reliable network community detection. *Physical Review E*, 100, 052308. <https://doi.org/10.1103/PhysRevE.100.052308>
- Comiso, J. C. (2012). Large decadal decline of the Arctic multiyear ice cover. *Journal of Climate*, 25(4), 1176–1193. <https://doi.org/10.1175/JCLI-D-11-00113.1>
- Cowen, R. K., & Sponaugle, S. (2009). Larval dispersal and marine population connectivity. *Annual Review of Marine Science*, 1(1), 443–466. <https://doi.org/10.1146/annurev.marine.010908.163757>
- Danon, L., Díaz-Guilera, A., Duch, J., & Arenas, A. (2005). Comparing community structure identification. *Journal of Statistical Mechanics: Theory and Experiment*, 2005(09), P09008. <https://doi.org/10.1088/1742-5468/2005/09/P09008>
- Dee, D. P., Uppala, S. M., Simmons, A. J., Berrisford, P., Poli, P., Kobayashi, S., et al. (2011). The ERA-Interim reanalysis: Configuration and performance of the data assimilation system. *Quarterly Journal of the Royal Meteorological Society*, 137(656), 553–597. <https://doi.org/10.1002/qj.828>
- Delandmeter, P., & van Sebille, E. (2019). The parcels v2.0 Lagrangian framework: New field interpolation schemes. *Geoscientific Model Development Discussions*, 12(8), 1–24. <https://doi.org/10.5194/gmd-2018-339>
- Donner, R. V., Lindner, M., Tupikina, L., & Molkenhain, N. (2019). Characterizing flows by complex network methods. In E. Macau (Ed.), *A mathematical modeling approach from nonlinear dynamics to complex systems* (pp. 197–226). Springer.
- Dubois, M., Rossi, V., Ser-Giacomi, E., Arnaud-Haond, S., López, C., & Hernández-García, E. (2016). Linking basin-scale connectivity, oceanography and population dynamics for the conservation and management of marine ecosystems: Large-scale connectivity and management of marine ecosystems. *Global Ecology and Biogeography*, 25(5), 503–515. <https://doi.org/10.1111/geb.12431>
- Edler, D., Bohlin, L., & Rosvall, a. (2017). Mapping higher-order network flows in memory and multilayer networks with Infomap. *Algorithms*, 10(4), 112. <https://doi.org/10.3390/a10040112>
- Fernandez, E., & Lellouche, J. (2018). Product user manual for the global ocean physical reanalysis product GLOBAL_REANALYSIS_PHY_001_030. Copernicus Marine Environment Monitoring Service.
- Fortunato, S. (2010). Community detection in graphs. *Physics Reports*, 486(3–5), 75–174. <https://doi.org/10.1016/j.physrep.2009.11.002>
- Fortunato, S., & Barthelemy, M. (2007). Resolution limit in community detection. *Proceedings of the National Academy of Sciences*, 104(1), 36–41. <https://doi.org/10.1073/pnas.0605965104>
- Froyland, G., Stuart, R. M., & van Sebille, E. (2014). How well-connected is the surface of the global ocean? *Chaos: An Interdisciplinary Journal of Nonlinear Science*, 24(3), 033126. <https://doi.org/10.1063/1.4892530>
- Good, B. H., De Montjoye, Y.-A., & Clauset, A. (2010). Performance of modularity maximization in practical contexts. *Physical Review E*, 81(4), 046106. <https://doi.org/10.1103/PhysRevE.81.046106>
- Goosse, H., & Fichefet, T. (1999). Importance of ice-ocean interactions for the global ocean circulation: A model study. *Journal of Geophysical Research: Oceans*, 104(C10), 23337–23355. <https://doi.org/10.1029/1999JC900215>
- Hadjighasem, A., Farazmand, M., Blazevski, D., Froyland, G., & Haller, G. (2017). A critical comparison of Lagrangian methods for coherent structure detection. *Chaos*, 27(5), 1–25. <https://doi.org/10.1063/1.4982720>
- Hadjighasem, A., Karrasch, D., Teramoto, H., & Haller, G. (2016). Spectral-clustering approach to Lagrangian vortex detection. *Physical Review E*, 93(6), 063107. <https://doi.org/10.1103/PhysRevE.93.063107>
- Haller, G., & Yuan, G. (2000). Lagrangian coherent structures and mixing in two-dimensional turbulence. *Physica D: Nonlinear Phenomena*, 147(3–4), 352–370. [https://doi.org/10.1016/S0167-2789\(00\)00142-1](https://doi.org/10.1016/S0167-2789(00)00142-1)
- Hansen, C., Kvaleberg, E., & Samuelsen, A. (2010). Anticyclonic eddies in the Norwegian Sea; their generation, evolution and impact on primary production. *Deep Sea Research Part I: Oceanographic Research Papers*, 57(9), 1079–1091. <https://doi.org/10.1016/j.dsr.2010.05.013>
- Hardesty, B. D., Harari, J., Isobe, A., Lebreton, L., Maximenko, N., Potemra, J., et al. (2017). Using numerical model simulations to improve the understanding of micro-plastic distribution and pathways in the marine environment. *Frontiers in Marine Science*, 4, 30. <https://doi.org/10.3389/fmars.2017.00030>
- Huffman, D. A. (1952). A method for the construction of minimum-redundancy codes. *Proceedings of the IRE*, 40(9), 1098–1101.
- Jacobi, M. N., André, C., Döös, K., & Jonsson, P. R. (2012). Identification of subpopulations from connectivity matrices. *Ecography*, 35(11), 1004–1016. <https://doi.org/10.1111/j.1600-0587.2012.07281.x>
- Kawamoto, T., & Rosvall, M. (2015). Estimating the resolution limit of the map equation in community detection. *Physical Review E*, 91(1), 012809. <https://doi.org/10.1103/PhysRevE.91.012809>
- Kheirkhazadeh, M., Lancichinetti, A., & Rosvall, M. (2016). Efficient community detection of network flows for varying Markov times and bipartite networks. *Physical Review E*, 93(3), 032309. <https://doi.org/10.1103/PhysRevE.93.032309>
- Lambiotte, R., & Rosvall, M. (2012). Ranking and clustering of nodes in networks with smart teleportation. *Physical Review E*, 85(5), 056107. <https://doi.org/10.1103/PhysRevE.85.056107>
- Lancichinetti, A., & Fortunato, S. (2012). Consensus clustering in complex networks. *Scientific Reports*, 2, 336. <https://doi.org/10.1038/srep00336>
- Logemann, K., & Harms, I. (2006). High resolution modelling of the North Icelandic Irminger current (NIIC). *Ocean Science*, 2(2), 291–304. <https://doi.org/10.5194/os-2-291-2006>
- Maximenko, N., Hafner, J., & Nilner, P. (2012). Pathways of marine debris derived from trajectories of Lagrangian drifters. *Marine Pollution Bulletin*, 65(1–3), 51–62. <https://doi.org/10.1016/j.marpolbul.2011.04.016>
- Meredith, M., Sommerkorn, M., Cassotta, S., Derksen, C., Ekaykin, A., Hollowed, A., et al. (2019). Polar regions [Book section]. In H. Pörtner, D. Roberts, V. Masson-Delmotte, P. Zhai, M. Tignor, & E. Poloczanska (Eds.), *IPCC special report on the ocean and cryosphere in a changing climate* (p. 203–320). Geneva, Switzerland: IPCC.
- Molkenhain, N., Rehfeld, K., Marwan, N., & Kurths, J. (2015). Networks from flows - from dynamics to topology. *Scientific Reports*, 4(1), 4119. <https://doi.org/10.1038/srep04119>
- Mysak, L. A., & Schott, F. (1977). Evidence for baroclinic instability of the Norwegian Current. *Journal of Geophysical Research*, 82(15), 2087–2095. <https://doi.org/10.1029/JC082i015p02087>
- NEMO System Team. (n.d.). *Nemo ocean engine (Computer software manual No. 27)*. Zenodo. <https://doi.org/10.5281/zenodo.1464816>
- Newman, M. E. J. (2010). *Networks*. Oxford University Press.
- Newman, M., & Girvan, M. (2004). Finding and evaluating community structure in networks. *Physical Review E*, 69(2), 026113. <https://doi.org/10.1103/PhysRevE.69.026113>
- Nurser, A. J. G., & Bacon, S. (2014). The Rossby radius in the Arctic Ocean. *Ocean Science*, 10(6), 967–975. <https://doi.org/10.5194/os-10-967-2014>

Padberg-Gehle, K., & Schneide, C. (2017). Network-based study of Lagrangian transport and mixing. *Nonlinear Processes in Geophysics*, 24(4), 661–671. <https://doi.org/10.5194/npg-24-661-2017>

Polyakov, I. V., Bhatt, U. S., Walsh, J. E., Abrahamsen, E. P., Pnyushkov, A. V., & Wassmann, P. F. (2013). Recent oceanic changes in the Arctic in the context of long-term observations. *Ecological Applications*, 23(8), 1745–1764. <https://doi.org/10.1890/11-0902.1>

Rossi, V., Ser-Giacomi, E., López, C., & Hernández-García, E. (2014). Hydrodynamic provinces and oceanic connectivity from a transport network help designing marine reserves. *Geophysical Research Letters*, 41(8), 2883–2891. <https://doi.org/10.1002/2014GL059540>

Rosvall, M., Axelsson, D., & Bergstrom, C. T. (2009). The map equation. *The European Physical Journal - Special Topics*, 178(1), 13–23. <https://doi.org/10.1140/epjst/e2010-01179-1>

Rosvall, M., & Bergstrom, C. T. (2011). Multilevel compression of random walks on networks reveals hierarchical organization in large integrated systems. *PLoS One*, 6(4), e18209. <https://doi.org/10.1371/journal.pone.0018209>

Rosvall, M., Delvenne, J.-C., Schaub, M. T., & Lambiotte, R. (2017). Different approaches to community detection. arXiv [physics]. <https://doi.org/10.1002/9781119483298.ch4>

Schaub, M. T., Delvenne, J.-C., Yaliraki, S. N., & Barahona, M. (2012). Markov dynamics as a zooming lens for multiscale community detection: Non clique-like communities and the field-of-view limit. *PLoS One*, 7(2), e32210. <https://doi.org/10.1371/journal.pone.0032210>

Schaub, M. T., Lambiotte, R., & Barahona, M. (2012). Encoding dynamics for multiscale community detection: Markov time sweeping for the map equation. *Physical Review E*, 86(2), 026112. <https://doi.org/10.1103/PhysRevE.86.026112>

Ser-Giacomi, E., Rossi, V., López, C., & Hernández-García, E. (2015). Flow networks: A characterization of geophysical fluid transport. *Chaos: An Interdisciplinary Journal of Nonlinear Science*, 25(3), 036404. <https://doi.org/10.1063/1.4908231>

Ser-Giacomi, E., Vasile, R., Hernández-García, E., & López, C. (2015). Most probable paths in temporal weighted networks: An application to ocean transport. *Physical Review E*, 92(1), 012818. <https://doi.org/10.1103/PhysRevE.92.012818>

Shannon, C. E. (1948). A mathematical theory of communication. *Bell System Technical Journal*, 27(3), 379–423. <https://doi.org/10.1002/j.1538-7305.1948.tb01338.x>

Strehl, A., & Ghosh, J. (2003). Cluster ensembles—A knowledge reuse framework for combining multiple partitions. *Journal of Machine Learning Research*, 3, 583–617. <https://doi.org/10.1162/15324430321897735>

Thomas, C. J., Lambrechts, J., Wolanski, E., Traag, V. A., Blondel, V. D., Deleersnijder, E., et al. (2014). Numerical modelling and graph theory tools to study ecological connectivity in the Great Barrier Reef. *Ecological Modelling*, 272, 160–174. <https://doi.org/10.1016/j.ecolmodel.2013.10.002>

van Sebille, E., Griffies, S. M., Abernathy, R., Adams, T. P., Berloff, P., Biastoch, A., et al. (2018). Lagrangian ocean analysis: Fundamentals and practices. *Ocean Modelling*, 121, 49–75. <https://doi.org/10.1016/j.ocemod.2017.11.008>

Vaughan, D., Comiso, J., Allison, L., Carrasco, J., Kaser, G., Kwok, R., et al. (2013). Observations: Cryosphere [Book section]. In T. F. Stocker, et al. (Eds.), *Climate change 2013: The physical science basis. Contribution of working group I to the fifth assessment report of the intergovernmental panel on climate change* (pp. 317–382). Cambridge, UK and New York, NY, USA: Cambridge University Press. <https://doi.org/10.1017/CBO9781107415324.012>

Wichmann, D., Kehl, C., Dijkstra, H. A., & van Sebille, E. (2020). Detecting flow features in scarce trajectory data using networks derived from symbolic itineraries: An application to surface drifters in the North Atlantic. *Nonlinear Processes in Geophysics Discussions*, 27(4), 501–518. <https://doi.org/10.5194/npg-2020-18>

References From the Supporting Information

Bohlin, L., Edler, D., Lancichinetti, A., & Rosvall, M. (2014). Community detection and visualization of networks with the map equation framework. In Y. Ding, R. Rousseau, & D. Wolfram (Eds.), *Measuring scholarly impact* (pp. 3–34). Springer Cham. https://doi.org/10.1007/978-3-319-10377-8_1

Hopkins, T. S. (1991). The GIN Sea—A synthesis of its physical oceanography and literature review 1972–1985. *Earth-Science Reviews*, 30(3–4), 175–318. [https://doi.org/10.1016/0012-8252\(91\)90001-V](https://doi.org/10.1016/0012-8252(91)90001-V)

Jones, E., Oliphant, T., Peterson, P. (2001). *SciPy: Open source scientific tools for Python*. Retrieved from <http://www.scipy.org/>

Karrer, B., Levina, E., & Newman, M. E. J. (2008). Robustness of community structure in networks. *Physical Review E*, 77(4), 046119. <https://doi.org/10.1103/PhysRevE.77.046119>

Khatiwal, S., Visbeck, M., & Cane, M. A. (2005). Accelerated simulation of passive tracers in ocean circulation models. *Ocean Modelling*, 9(1), 51–69. <https://doi.org/10.1016/j.ocemod.2004.04.002>

Maneewongvatana, S., & Mount, D. M. (1999). It's okay to be skinny, if your friends are fat. In *4th annual CGC workshop on computational geometry* (p. 8).

McAdam, R., & van Sebille, E. (2018). Surface connectivity and Inter-ocean exchanges from Drifter-based transition matrices. *Journal of Geophysical Research: Oceans*, 123(1), 514–532. <https://doi.org/10.1002/2017JC013363>

Meilü, M. (2007). Comparing clusterings—an information based distance. *Journal of Multivariate Analysis*, 98(5), 873–895. <https://doi.org/10.1016/j.jmva.2006.11.013>

Moresi, L., & Mather, B. (2019). Stripy: A Python module for (constrained) triangulation in Cartesian coordinates and on a sphere. *Journal of Open Source Software*, 4(38), 1410. <https://doi.org/10.21105/joss.01410>

Regan, H. C., Lique, C., & Armitage, T. W. K. (2019). The Beaufort Gyre extent, shape, and location between 2003 and 2014 from satellite observations. *Journal of Geophysical Research: Oceans*, 124(2), 844–862. <https://doi.org/10.1029/2018JC014379>

Rosvall, M., & Bergstrom, C. T. (2008). Maps of random walks on complex networks reveal community structure. *Proceedings of the National Academy of Sciences*, 105(4), 1118–1123. <https://doi.org/10.1073/pnas.0706851105>

Sadourny, R., Arakawa, A., & Mintz, Y. (1968). Integration of the nondivergent barotropic vorticity equation with an icosahedral-hexagonal grid for the sphere. *Monthly Weather Review*, 96(6), 6.

Sætre, R. (1999). Features of the central Norwegian shelf circulation. *Continental Shelf Research*, 19(14), 1809–1831. [https://doi.org/10.1016/S0278-4343\(99\)00041-2](https://doi.org/10.1016/S0278-4343(99)00041-2)

van Sebille, E. (2014). Adrift.org.au — A free, quick and easy tool to quantitatively study planktonic surface drift in the global ocean. *Journal of Experimental Marine Biology and Ecology*, 461, 317–322. <https://doi.org/10.1016/j.jembe.2014.09.002>

Viamontes Esquivel, A., & Rosvall, M. (2011). Compression of flow can reveal overlapping-module organization in networks. *Physical Review X*, 1(2), 021025. <https://doi.org/10.1103/PhysRevX.1.021025>

Wang, N., & Lee, J.-L. (2011). Geometric properties of the icosahedral-hexagonal grid on the two-sphere. *SIAM Journal on Scientific Computing*, 33(5), 2536–2559. <https://doi.org/10.1137/090761355>