Studying structural change in the European Aviation Network

Juul Vossen
Psych
Vrije Universiteit Brussel
Elsene, Belgium
mailto:juul.vossen@vub.be

Enrico Spinielli
Rainer Koelle
Aviation Intelligence Unit
EUROCONTROL
Brussels, Belgium
enrico.spinielli@eurocontrol.int
rainer.koelle@eurocontrol.int

***Abstract*— Drastic loss of flight connections due to the COVID-19 pandemic has called for new approaches to accurately study structural change in the European Aviation Network. This study highlights the limitations of traditional centrality-based network approaches and proposes a diffusion-based graph embedding approach using the GraphWave algorithm. This new approach was validated using domain knowledge and tested in its ability to capture known events that occurred during and after the COVID-19 pandemic. The network is modelled based on all flights departing from and arriving to European airports in the period of 2019 through 2022. Flight connections were aggregated on a weekly basis to analyze structural embeddings and the structural role of airports. The temporal analysis supported the identification and assessment of changes to the role of airports and structural changes of the network. This study shows the potential of the approach by applying the model to uncover global, regional, and local change dynamics, and highlighting its potential as a valuable tool for researchers and practitioners studying the evolution of complex networks.**

Keywords-component; network analysis, graph representation learning, change dynamics

I. Introduction

The COVID-19 pandemic had a significant impact on the European Aviation Network, resulting in an unprecedented loss of flight connections. The travel restrictions imposed by governments to curb the spread of the virus severely limited the demand for air travel, leading to a dramatic decline in passenger traffic. Following the drop in 2020, traffic in the EUROCONTROL (ECTL) area recovered to 6.2 million flights in 2021. This corresponds to roughly half of the traffic in 2019 [1].

Differences in policies implemented by different countries to address the pandemic, such as the timing and severity of travel restrictions, as well as the effectiveness of their response to the crisis, made the effect of the pandemic on the aviation network highly variable [2,3,4]. Both within and between countries, airports differ in the extent to which they lost flights and connections to other airports. Departing from the notion that network structure is an emergent property of the structural roles of the individual airports that make up the network, this raises the question whether the network still functions the same way as it did before the pandemic, or whether these local changes in the network have led to fundamental changes in the network structure.

Answering this question may not be straightforward, as traditional centrality based network analysis techniques are ill-equipped to deal with both: (i) the strong changes in network activity and (ii), the accurate capture of structural changes in the network. First, to detect meaningful structural changes over the course of the pandemic, we need to be able to compare the network at any given time to a reference network (i.e., pre-covid network). This poses challenges, as traditional centrality measures are generally created for time-independent networks [5]. This means that extra care has to be taken when comparing traditional centrality measures across time as scaling issues may arise due to their direct dependence on the structure of the graph. Think for instance of the direct dependence of flight activity (i.e., edge weights) on centrality measures such as weighted degree or weighted betweenness centrality, or the difference in the magnitude of the dominant eigenvalue between different networks that plays an important role in eigenvector based centrality measures.

Aside from scaling issues, traditional centrality measures such as PageRank, betweenness and closeness also tend to have a poor conceptual fit with the notion of structural changes, as the importance of a node is always seen as a function of all the nodes in the network [6,7]. In other words, the centrality measure of a node is directly dependent both on the centrality of all the other nodes in the network, as well as the number of nodes that are present in the network. This is problematic when we are interested in looking at local structural changes in a network, given that the local structural role of a node should not be dependent on changes in distant parts of the network.

Taking these issues into account, to determine changes in the functioning of the network with respect to the pre-covid network, we need an approach that can accurately capture the role of each airport in its local network neighborhood, while simultaneously being able to accurately compare this structural role with a baseline in the pre-covid network.

To help fill this gap, we propose to use a combination of a state-of-the-art unsupervised graph representation learning algorithm (i.e., GraphWave [8]) and K-means clustering to (i) derive a structural embedding for each airport in the network that is robust to scaling issues and (ii) find clusters of airports that have a similar structural role within their local network topology. Specifically, we propose to aggregate the European Air Transportation Network from January 1st 2019 to December 31st 2022 into weekly snapshots consisting of undirected networks in which the airports are represented as nodes and the number of flights as the weights along the edges. We then create a unique weekly baseline model with the snapshots corresponding to the weeks in 2019 (i.e., pre-covid), which will then be used to classify the structural embeddings derived from the snapshots of the corresponding weeks in 2020, 2021, and 2022 (i.e., during- and post-covid).

II. Background

The GraphWave algorithm is a diffusion-based node embedding algorithm from the field of graph representation learning. Graph representation learning is a field of machine learning that focuses on learning meaningful and compact representations of graph-structured data. The goal is to map the nodes and edges of a graph to a low-dimensional vector space, such that the properties of the graph are preserved in the embedding, allowing for downstream tasks such as link prediction and node classification.

While most network embedding techniques tend to model the proximity between nodes in a network, there has been increasing attention towards structural embeddings which focus on identifying node equivalences. Such equivalences are collections of nodes that share a similar role in the network (e.g. a hub, or a bridge between parts of the network), irrespective of their location in the network [6]. In contrast to traditional proximity based approaches—which consider nodes in close proximity to one another to be more similar than nodes that are more distant—structure based approaches draw on the intuition that nodes that share a similar role in the network also perform a similar function in the network [5].

A clear example of such equivalences with regard to the European Aviation Network can for instance be seen when looking at airports such as Heathrow, Schiphol, and Istanbul Airport. While these airports are located in different regions of the network, they share a similar role in the network. All three airports serve as hubs, connecting large parts of the network both within, as well as outside of Europe.

GraphWave uses spectral graph wavelets to analyze the diffusion patterns of a heat kernel centered at each node. These patterns consist of wavelet coefficients, which are the amount of energy that is passed from the target node to each other node in the network before the signal has decayed. These wavelet coefficients are then treated as a probability distribution by sampling the empirical characteristic function for each coefficient. Intuitively, the algorithm passes a heat signal through each airport and detects how far it can reach in the network before cooling off. The more connections an airport has, the more places the signal can spread to. Likewise, the more flights on a given connection, the more energy can be passed along the route and the further the signal can reach before decaying. Notice that strong heat propagation is not only dependent on the number of flights and flight routes departing from the starting airport, but also on the amount of flights and connections of its neighbors. This means that airports with a similar number of flights and flight routes can have a vastly different role in the network depending on which airports they are connected to. Think for instance of the hypothetical case of two airports who only have one flight connecting them to another airport. The first airport is connected to an airport with connections to two other airports, while the second airport is connected to an airport which has connections to sixty other airports. While both starting airports are identical, the propagation of heat will be vastly different. Moreover, this also means that the role of an airport is not only dependent on changes in the number of flights and flight routes it has with its direct neighbors, but also on changes in airports that are connected to its neighbors.

III. Data and Methodology

A. Data

This study builds on all flown flights arriving or departing from an airport[1] within one of the EUROCONTROL (ECTL) member states between 01-01-2019 and 31-12-2022. See Table 1 for an overview of the total number of flights and airports connected to the network per year.

TABLE 1: TOTAL NUMBER OF FLIGHTS AND AIRPORTS PER YEAR

Year	Total # of Flights	Total # of Airports		
		Intra-ECTL	Non-ECTL	Total
2019	10.846.405	2295	1232	3527
2020	4.861.490	2242	1141	3383
2021	6.034.669	2311	1178	3489
2022	9.052.301	2327	1227	3554

The data are aggregated on a weekly level to ease computation, while still being able to capture seasonal variability [9]. Specifically, we create an undirected network for each week where the airports serve as the nodes of the network, the flight routes as the edges, and the number of flights as the edge weights.

To ensure reproducibility of the results, the script to run the analysis (including sample data) can be found on the following GitHub repository: https://github.com/euctrl-pru/aviation-network-structure-model

[1] Land airports with scheduled regional airline service, or regular general aviation or military traffic classified by the OurAirports database as "medium sized". This naming convention also labels major airports as "midsized".

B. *Analysis Plan*

Create structural embeddings. To create structural embeddings for each week, we use the adjacency matrices as input for the Python implementation of the GraphWave algorithm by the Stanford Network Analysis Project (SNAP)[10]. The default settings [8] of the algorithm are used to find the optimal scale values for the embeddings of the baseline weeks. As noted, the scale of the signal determines the radius of the network neighborhood around each node, where smaller scale values allow the signal little time to propagate across the network, while large scale values cause the heat to become equally spread across all the nodes of the network. The algorithm calculates an interval bounded by a minimum and a maximum scale value based on the analysis of variance of heat diffusion wavelets. This means that we accommodate for the fact that networks may have different ideal scale values depending on the seasonal activity. For the during- and post-covid weeks, we use the corresponding optimal scale values as the input for the scale of the signal, which ensures that the embedding for each week is created using the same signal as the corresponding baseline week.

Create and assess classification model validity. A K-means clustering model is trained for each of the baseline weeks. The centroids are used to classify the embeddings of the networks of the corresponding weeks in 2020, 2021, and 2022. This allows for the evaluation of each airport's heat signature as if it had occurred in the corresponding pre-covid week. To evaluate the model validity, different values of k for the K-means algorithm are evaluated to determine the ideal number of clusters using the Calinski-Harabasz (CH) Index. The CH Index measures the similarity for each object to its own cluster compared to the other clusters, where higher values indicate that clusters are dense and well separated [11]. The clustering for each of the weeks in the baseline year are compared to determine the consistency of the classification of our model.

Additionally, the robustness of the model will be tested by running the same analysis on reduced versions of the network. The goal of this analysis will be to test the extent to which the local embeddings are dependent on the overall structure of the network, and assess the models ability to handle loss of connectivity. The first reduction will remove the 50% airports with lowest activity and will test the influence (the loss) of smaller airports on the embeddings of the top 50% airports. The second reduction will remove all weights and directionality (i.e., incoming and outgoing flights) and only look at the presence of flight routes between airport pairs.

Validate predictions using domain knowledge. Domain knowledge is used to cross-reference known events in the period of 2020-2022 to assess the validity of our predictions using the centroids of the baseline model. Specifically, four known events that have occurred in this period are used: (i) the grounding of unused aircrafts during the pandemic, (ii) the stability (and slight increase) of cargo flights during the pandemic, (iii) the large-scale airline strikes in Belgium and Germany in 2022, and (iv) the closure of the Ukrainian Airspace to all civil traffic on the 24th of February 2022.

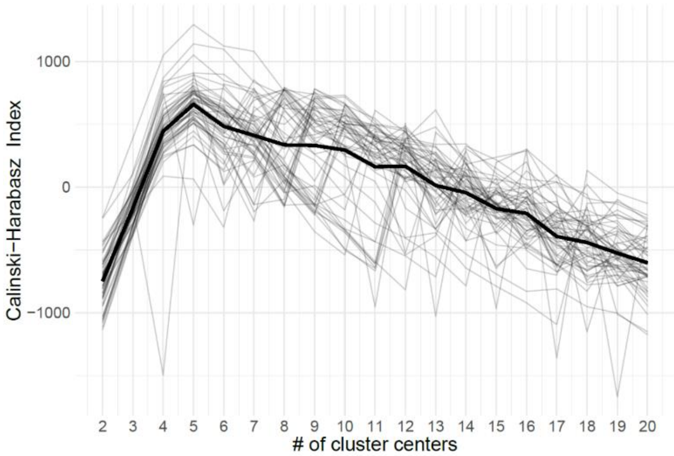

Figure 1: Calinski Harabasz index scores centered by week for models ranging from 2 to 20 centers

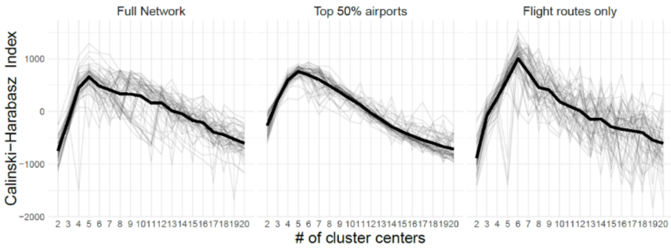

Figure 2: Calinski Harabasz index scores for the full network, top 50% active airports, and flight routes only network.

IV. Results

A. *Model Validation*

Plotting the CH index[2] for all 53 weeks shows that, on average, there is a peak around five centers, but that there is a considerable amount of variability around the nine centers region (c.f. Figure 1). Ranking the performance of each choice of centers based on their CH-index and calculating the mean rank score for the fifty-three models, shows that on average the models with five centers rank the highest, followed by six centers, seven centers and ten centers. This means that the difference between the signals varies in intricacy across the weeks of the year. There is a trade-off to be made between choosing the number of centers that either overfits or underfits the data. Having too many centers runs the risk of finding patterns that are only stochastic noise, while having too few centers may cause distinct patterns to be lumped together, causing them to be obscured from our analysis. This analysis is based on the model using six centers as it allows for a more expressive model than the overall best solution (i.e., five centers), while still ranking as the second-best number of clusters.

[2] The values of the CH-index are centered by subtracting the mean CH-index value for each week

TABLE 2: CLUSTER DESCRIPTIVE STATISTICS: MEAN AND STANDARD DEVIATION IN WEEKLY NUMBER OF CONNECTED AIRPORTS, FLIGHT ROUTES, AND FLIGHTS FOR EACH CLUSTER IN 2019

Cluster	*# of Airports in cluster*	*Mean # Airports Connected*	*SD # Airports Connected*	*Mean # Flight routes*	*SD # Flight routes*	*Mean # Flights*	*SD # Flights*
C1	10289	4	2	5	3	9	8
C2	13313	13	8	18	12	50	35
C3	8847	33	18	50	28	167	91
C4	5028	77	33	128	53	594	285
C5	2936	142	47	248	77	1761	610
C6	1526	229	57	419	103	5678	2104

Testing the model on the reduced networks, it is interesting to see that the optimal scales values as calculated by the algorithm are identical for the full network and the network containing only the presence of flight routes ($\overline{\text{range}}$ = [1.551; 6.746]). This indicates that it is not the amount of activity in the network that determines the optimal signal value, but rather the underlying structure of the network. This is similarly reflected in the smaller network containing the top 50% airports, where we find a smaller average optimal scale range ($\overline{\text{range}}$ = [1.314; 5.716]) compared to the full network.

In terms of CH-index (c.f. Figure 2) we see that the signal of the top 50% airports network is more stable compared to the full network, with a clearer preference ranging around six centers. With respect to the flight routes only network, we see a noisy signal similar to the full network. However, in contrast to the full network, we find a more pronounced preference for a six centers solution.

With respect to the node embeddings (i.e., heat signals), and the derived clusters, we find overall very strong correlations between the full network and both reduced networks. For the top 50% only network, we find an almost perfect correlation with respect to the node embedding ($r_{embedding}$(6968898) = .99, p <.001) with an average squared distance of .0001 for each heat signal. In terms of the derived clusters we find a similarly strong correlation for the clustering derived using six centers (i.e., the amount of centers used in our working model) ($r_{clustering}$(69687) = .96, p <.001).

For the flight routes only network, we again find an almost perfect correlation with respect to the node embeddings ($r_{embedding}$(9710498) = .99, p <.001), with an average squared distance of .0003 per heat signal. With respect to the clustering using six centers we also find a slightly lower but strong positive correlation ($r_{clustering}$(97103) = .88, p <.001).

Combing the findings, it appears that the node embeddings derived by the GraphWave algorithm are robust to various degree of loss of connectivity. Reducing the entire graph to its bare structure (i.e., flight routes only), the model is still able to capture meaningful local structural roles comparable to the full network, making the model a good fit to study structural changes with respect to the loss of connectivity caused by the pandemic.

B. Cluster descriptive statistics

To get a feel for the structural roles within the network, we provide the average and the standard deviation of the number of connected airports, flight routes and flights for each cluster, as well as the amount of times an airport is classified as being in the respective cluster across the weeks of 2019 (see Table 2). Figure 3 shows normalized heat propagation for exemplars for each cluster

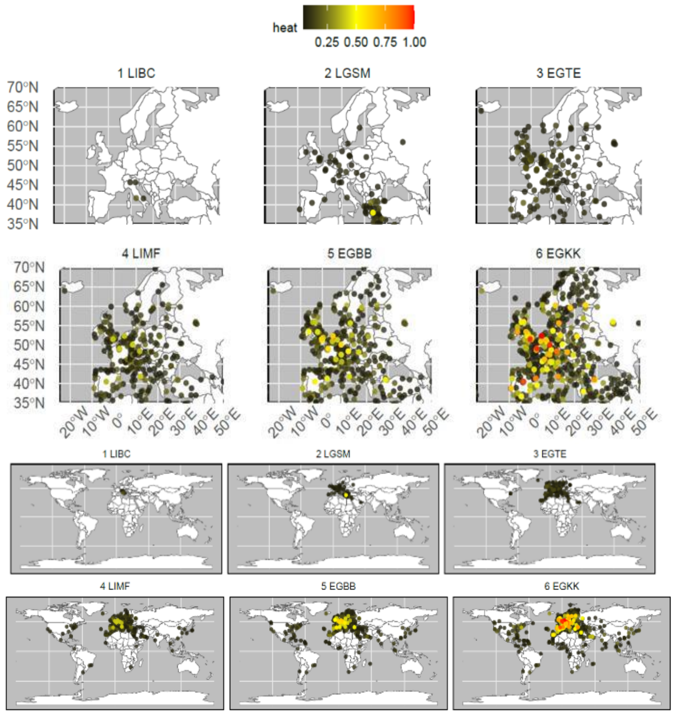

Figure 3: Heat propagation for Crotone Airport (LIBC), Samos International Airport (LGSM), Exeter International Airport (EGTE), Turin Airport (LIMF), Birmingham Airport (EGBB), and London Gatwick Airport (EGKK), which represent cluster C1 to C6 respectively. The amount of heat that is propagated by each airport is normalized across the clusters to show differences in propagation between clusters.

During 2019, the majority of the roles within the EATN fall within cluster C2 (N = 13313), followed by cluster C3 (N = 8847) and cluster C1 (N = 10289). This means that over 50% of the airports in the network are on average connected to less than 13 airports, with less than 18 flight routes and less than 50 flights in a week.

Both the overall number, as well as the variability in the number of airports connected, flight routes, and flights increases for each increase in cluster. This denotes that airports are more similar in the lower clusters compared to the higher clusters. Moreover, this indicates that for higher clusters more emphasis is placed on the number of connected airports and flight routes,

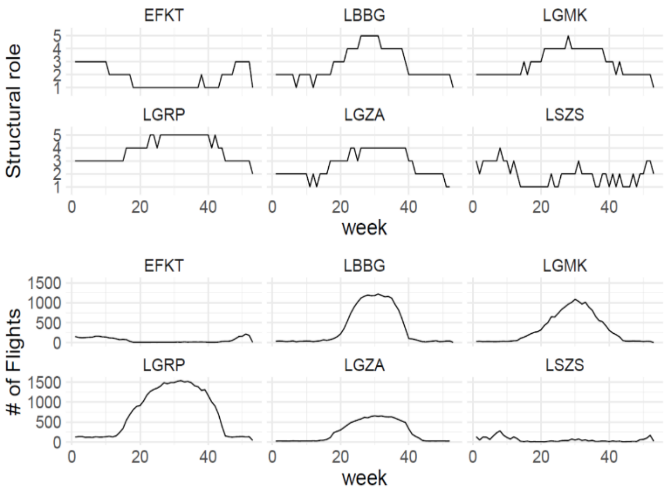

Figure 4: Six airports with highest variability in structural role in 2019. Note: EFKT = Kittilä Airport; LBBG = Burgas Airport; LGMK = Mykonos Airport; LGRP = Rhodos International Airport; LGZA = Zakynthos International Airport; LSZS = Samedan Airport.

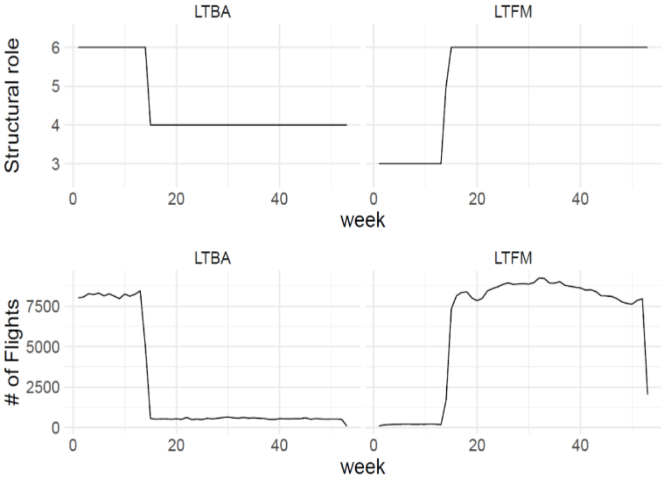

Figure 5: Evolution of structural role of Atatürk Airport (LTBA) and Istanbul Airport (LFTM) in 2019.

rather than the absolute number of flights, indicating that it is not only the amount of flights that is the driving factor for a high structural role, but also the level of connectivity of the connected airports.

C. *Classification consistency*

To verify the consistency of our model, we compute a Pearson correlation coefficient for each pair of weekly classifications. We find a high overall correlation between the weeks ($\bar{r}_{overall}$ = .92, p <.001), with high correlations between consecutive weeks ($\bar{r}_{consecutive}$= .96, p <.001), and lower correlations between holiday-season and off-season week pairs ($r_{smallest}$= .82, p <.001). Overall, it appears the strength of the correlations follow a seasonal pattern [15]. Airports that show the highest variability in their role in the network are airports that handle mostly seasonal traffic. Figure 4 shows the overlap between the increase in flight activity and the increase structural role for Burgas Airport in Bulgaria, Mykonos Airport, Rhodes International Airport, and Zakynthos International Airport in Greece, Samedan Airport in Switzerland, and Kittilä Airport in

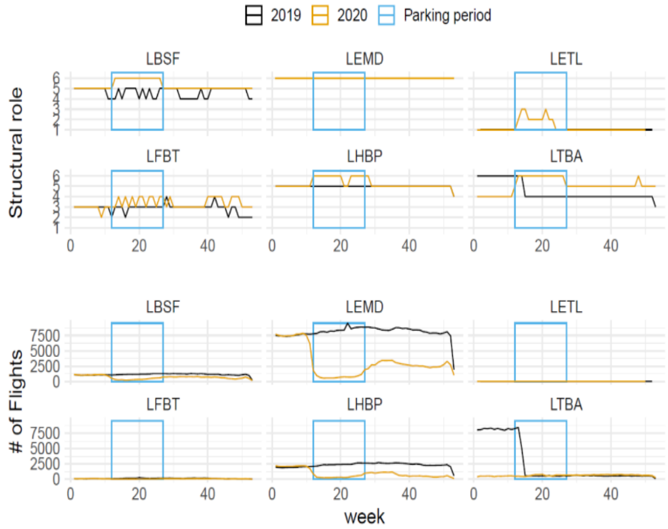

Figure 6: Evolution of structural role and total number of flights of the six airports that stored the largest amount of inactive aircrafts in 2020. Note: LBSF = Sofia Airport; LEMD = Adolfo Suárez Madrid–Barajas Airport; LETL = Teruel Airport; LFBT = Aéroport de Tarbes-Lourdes-Pyrénées; LHBP = Budapest Ferenc Liszt International Airport; LTBA = Atatürk Airport

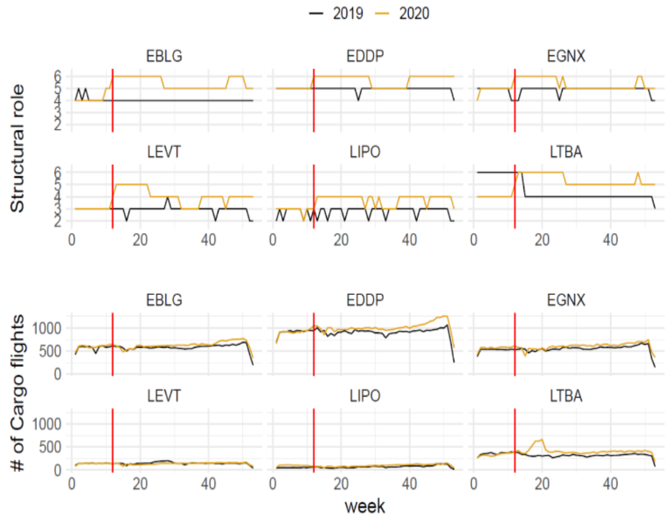

Figure 7: Evolution of structural role and total number of cargo flights of six airports that served the largest amount of cargo flights in 2020. Note: EBLG = Liège Airport; EDDP = Leipzig/Halle Airport; EGNX = East Midlands Airport; LEVT = Vitoria Airport, LIPO = Aeroporto di Brescia; LTBA = Atatürk

Finland. Note that while Kittilä Airport only has an average of 66 flights per week (with a max of 216 flights during its peak season), it shows a relatively strong leap in connectivity during its busiest season. This is a perfect example to illustrate the effect of being connected to well-connected neighbors, as during its seasonal peaks Kittilä becomes connected to some of the most connected airports in the European Aviation Network, such as Schiphol, Paris-Charles de Gaulle, London-Gatwick to name a few.

Interestingly, within-year role variability also picks up important local changes. For instance, Figure 5 shows the model output for the transfer of all scheduled commercial passenger flights from Atatürk Airport to Istanbul Airport on 6

April 2019. This is present in the structural embedding, as Istanbul Airport's cluster membership varied between three, four, and five for the first 14 weeks of 2019, but as of week 15 (April-8 – April-14) has risen to cluster six, at which it has stayed for the remainder of 2019. Conversely, Atatürk Airport is classified as cluster six for the first 14 weeks of 2019, after which it consistently dropped to cluster four, with three times a rise to cluster five in week 31, 45, and 50.

D. Validating predictions using domain knowledge

Given that we have no ground truth labels to test our model predictions, we make use of domain knowledge to assess the validity of the predictions. Starting with the events during the pandemic, we first look at the grounding of aircraft by airline operators due to the drastic reduction in (air) traffic demand. The idea is that the grounding of aircraft creates new "artificial" connections between airports that would otherwise not be connected. This causes the local networks of the origin airport and the destination to become weakly connected. This weak connection causes the signal to propagate too much more places across the network, causing airports connected by this link to get an increase in their structural role.

Based on the analysis of the Performance Review Unit (PRU) of EUROCONTROL[12] , we plot the evolution of the structural role of the six airports with the largest amount of inactive aircrafts (c.f. Figure 6). Looking at the time period when the majority of groundings took place, we see that all the airports experienced an increase in their structural role. A noteworthy exception is Adolfo Suárez Madrid–Barajas Airport which is consistently in the highest structural role.

Aside from an increase of inactive aircrafts, the pandemic resulted in an increase in the number of cargo flights operated within the European network. Looking at the six airports with the highest number of cargo flights, and whose market segment consist primarily of cargo flights (i.e., greater than 60%) , we see that this increase is reflected in their structural role in the network. All six airports see an increase in their structural role over the course of the pandemic (c.f. Figure 7). Specifically, we see a strong increase for Liège Airport and Atatürk Airport resulting in a cluster change from cluster 4 to cluster 6. Vitoria Airport in Spain makes a similar jump, jumping from cluster 3 to cluster 5.

Looking beyond 2020 to the aftermath of the pandemic, we see an increase in strikes among airline employees as the industry continues to struggle financially. The decrease in travel demand caused by the pandemic led to significant financial losses for airlines, and as a result, companies were forced to make cutbacks such as reducing routes and laying off employees. These cutbacks triggered tensions between the airlines and their employees resulting in strike actions to protest reduced pay and benefits. For example, strikes of Lufthansa's ground staff at Frankfurt Airport in Germany in July of 2022 prompted more than 600 flights to be canceled. Likewise, a national strike in Belgium in November of 2022 caused Brussel-Zaventem to cancel up to half of its flights, while Brussels Charleroi cancelled all Ryanair flights.

While having a profound impact on the daily traffic flow, none of the strikes appear to have shocked the network severely enough to result in a change in its structural role. One main reason for this is that our model is trained on weekly data. Accordingly, the effect of the loss of momentary connectivity is absorbed by the remaining days in the week. In other words, only prolonged events will be picked up by the model.

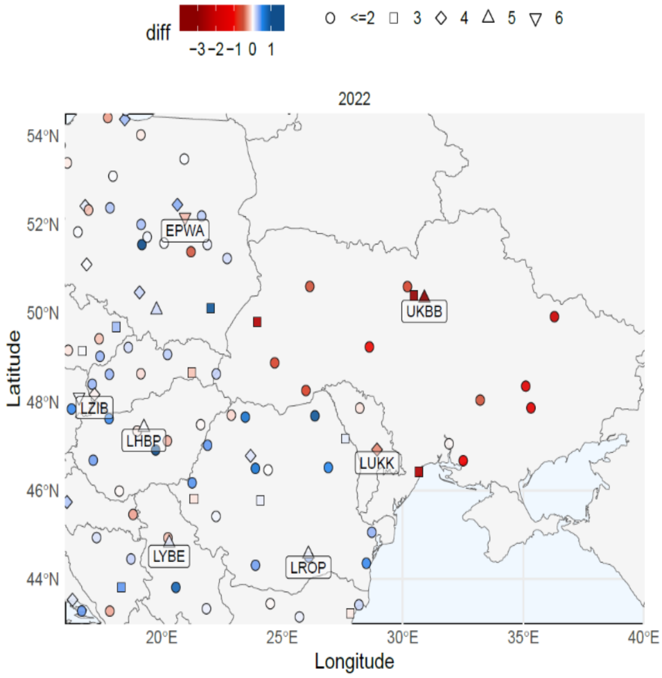

Figure 8: Change in structural role of airports in the Ukraine region in 2022. Note: Shape denotes average structural role in 2019. Color denotes difference of current average structural role compared to 2019. EPWA = Warsaw Chopin Airport; LROP = Henri Coandă International Airport; LYBE = Belgrade Nikola Tesla Airport; LZIB = Bratislava Airport; LUKK = Chișinău International Airport; UKBB = Boryspil International Airport; LHBP = Budapest Ferenc Liszt International Airport

One such event is the closure of the Ukrainian airspace to all civil traffic on the 24th of February in 2022. Figure 8 shows the change in structural roles for the airports in the Ukraine region. Unsurprisingly, all airports in Ukraine have drastically reduced in their structural role with the exception of one airport close to the border with Hungary. In terms of the structural role of the airports in the neighboring regions, we only find a decrease for airports in Moldova. Rather counter-intuitively, we find an increase in the structural role for the majority of the smaller airports in Romania, Hungary, Poland, Serbia, and Slovakia. The structural role of larger airports such as Henri Coandă International Airport, Belgrade Nikola Tesla Airport, and Bratislava Airport appears to have remained largely unchanged, except for Warsaw Chopin Airport, which has slightly decreased in structural role.

E. Change in role dynamics

In the following section, we will show how the model can guide us to uncover interesting change dynamics that might have been difficult to find using traditional approaches.

On a global level, looking at the different distributions of average cluster membership for the pre-, during-, and post-covid years, Figure 9 shows an interesting pattern. The pandemic seems to have caused a shift in the connectivity of the airports where we no longer see the strong equilibria around consistent classification (i.e. peaks around 1, 2, 3, 4, 5, 6), but rather a pattern of diffusion filling the gaps between the peaks. This indicates that airports that had consistent roles in the pre-covid network experienced a shift in connectivity due to changes in the network. There is an increase in connectivity for airports in the lower connected roles, and a slight decrease in the airports that were strongly connected in 2020, with a gradual recovery to the equilibria of the pre-covid network in 2022. Moreover, it seems that the overall connectiveness in the network in 2022 (M = 2.65, SD = 1.35) has slightly increased compared to 2019 (M = 2.56, SD = 1.35), $t(41445)$ = 25.07, $p < .001$.

Next, filtering by region, we find that, compared to 2019, the Ukraine is the region with both the highest overall cluster mobility, as well as the highest decrease in cluster membership. Looking at the United Kingdom—the region with the second greatest overall decrease in connectivity (and third in the region with the highest overall cluster mobility behind Ukraine and France)—we see a mixed impact of the pandemic on the connectivity of the UK airports (c.f. Figure 10).

Especially in the London area, a notable difference can be seen, both during and after the pandemic. For instance, the central role of Heathrow in the European network has remained unchanged throughout the pandemic, while Gatwick and Manchester airport both dropped a cluster over the course of the pandemic and recovered in 2022 (i.e. returned to their previous cluster). However, among the London airports: Heathrow, Gatwick, London Stansted, Luton, and London City, two stand out; Luton and London City.

Luton is the only airport among the five that went up in connectivity during the pandemic and still has higher connectivity compared to 2019. Conversely, London City is the only airport that has not recovered from the pandemic (c.f. Figure 11).

V. Discussion

The results of this study showcase the potential of using GraphWave to analyze and assess the evolution of the European Aviation Network by studying the changes in the structure of the network.

A. Model validity

Overall, we find that the structural roles derived by the model are robust to various levels of connectivity loss, showing strong correlations between the structural roles derived from the full network and the roles derived from the reduced networks. In terms of consistency, we find strong correlations between the clustering of consecutive weeks, as well as slightly weaker, but still strong correlations between weeks differing in season, which is in line with the expected seasonal variability [9]. Moreover, we find that the role mobility of an airport strongly relates to seasonal activity for the airports with the highest variability in structural role. The model also accurately picks up persistent local changes as seen with (i) the transfer of commercial flight from Atatürk Airport to Istanbul Airport, (ii) the increase of connectivity of airports that stored inactive aircrafts, (iii) and the increase of connectivity of majority cargo oriented airports in 2020. The

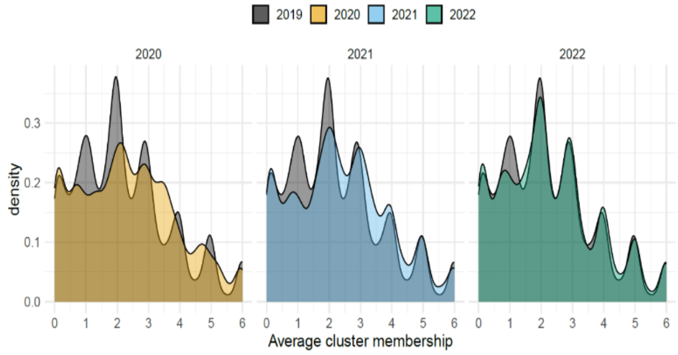

Figure 9: Distribution of average cluster membership from 2019-2022

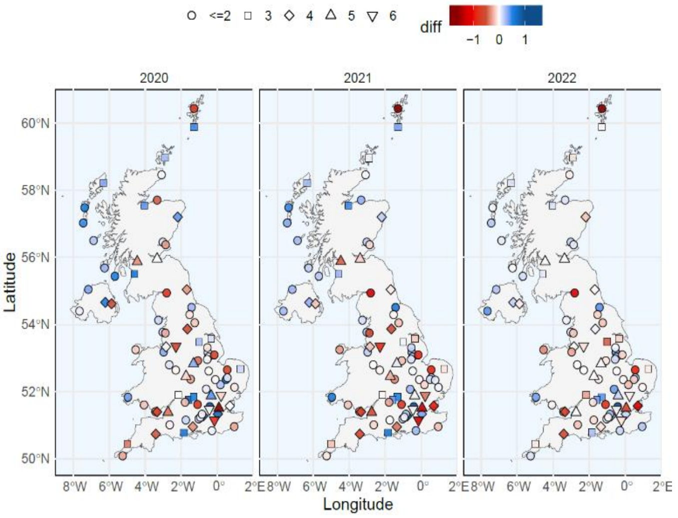

Figure 10: Evolution of structural roles of airports in the United Kingdom with respect to 2019. Note: Shape denotes average structural role in 2019. Color denotes difference of current average structural role compared to 2019.

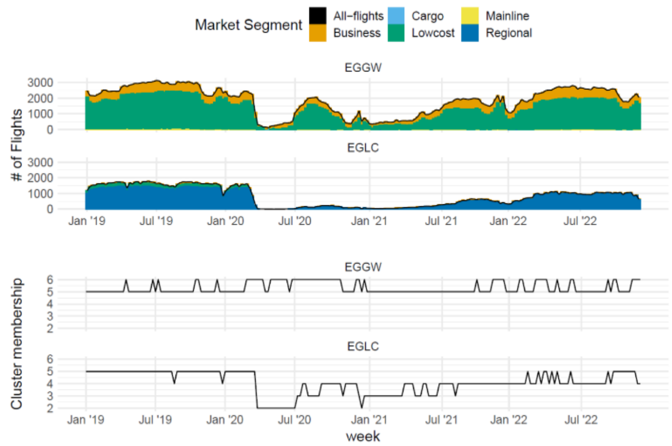

Figure 11: Evolution of market segment coverage and cluster membership for London Luton (EGGW) and London City (EGLC) Airport

SESAR Innovation Days
27-30 November 2023
Seville

model is not able to pick up the different strikes in 2022, nor does it find strong structural changes in the countries neighboring Ukraine. The reason for this may lie in the fact that the model is trained using aggregated weekly data which alleviates the impact of the single day strikes. Second, the data used for this study are origin-destination pairs of airports . In other words, they do not take into account flight routes, which may give a biased representation of the effect of the war on the structural roles in the neighboring countries.

B. *Model application*

In terms of model application, we find some interesting patterns regarding of the evolution of the structure of the network. On a global level, the COVID-19 pandemic has caused a shift in the connectivity of airports. Rather surprisingly, there appears to be a diffusion of connectivity filling the gaps between previously consistent classifications. This indicates that airports that had consistent roles in the pre-COVID-19 network experienced a shift in connectivity during the pandemic due to changes in the network. Specifically, it seems that the lower connected roles gained in connectivity, while airports that were previously strongly connected suffered a decrease in connectivity. Moving past COVID-19 we see a gradual recovery to the equilibria of the pre-COVID-19 network in 2022, with a slight increase in connectivity compared to 2019.

This paints an interesting, but confusing picture, as it feels counter intuitive that connectivity would increase in a time of an unprecedented drop in flight activity. While further research is needed to fully understand this phenomenon, it may be important to rethink what we mean with the concept of connectivity—especially in times with large restrictions on travel. While we normally speak of connectivity in terms of the time it takes for people to reach any destination in the network from a given airport, connectivity during times of travel restrictions can better be viewed as the potential for connectivity. In other words, less connected airports in the pre-covid had a greater potential for connectivity during the pandemic given that there would not have been any travel restrictions.

One possible explanation for instance may be that regional airports may have lost single connections to larger hubs, but in return gained more connections with other regional airports, thereby strengthening their role in a given local structure. Another possible explanation for the increase in connectivity may be due to the grounding of inactive aircrafts. Many planes were stored in regional airports. Transferring aircrafts from larger airports to these smaller regional airports may have caused an increase in the structural role of the airports within their neighborhoods, as they became artificially connected to one another due to the transfer of the aircraft. Future research can therefore test how the connectivity is affected by these artificial connections.

On a regional level, we find that, apart from Ukraine, overall, the airports in the United Kingdom underwent the strongest decrease in their structural role compared to 2019. There are considerable differences in the extent to which airports have endured and recovered from the effects of the pandemic. Specifically, zooming in on two London based airports with opposite change trajectories (i.e., Luton Airport and London City Airport), we find a clear difference in the market segment that both airports operate in. Future research can investigate the effect of market segment on the recovery rate, and if this effect is equal among all sizes of airports.

C. *Limitations of the model*

In terms of limitations, we find that one of the main limitations of the model is that it has a strong dependency on the available data. In other words, the accuracy and reliability of the models' classifications is directly dependent on the airports and flights that are included in the modelled network. For example, initially, the study only included flights arriving at and departing from airports starting with ICAO code L and E (i.e. Northern Europe, and Southern Europe). The model then classified Manchester Airport in cluster five, despite being the third busiest airport of the United Kingdom in 2019. Adding all flights arriving and departing from EUROCONTROL member states (including Morocco, Ukraine, and Iceland) , caused Manchester Airport to be classified as cluster 6. This indicates that it is important to be aware of which airports and flights are present in the data (and which are not), and that any inferences made should be framed within the data that is used.

Another limitation of the approach is that it relies on an unsupervised learning algorithm to extract the structural embeddings and identify the structural roles in the network. While also a strength, given that we can run the analysis on any network without needing pre-classified airports, not having ground truth labels means that it is important to have access to domain knowledge before we can confidently draw inferences from the model.

A last limitation that should be discussed is the effect of clustering on the sensitivity of the model. Given that we are working with clusters, we should be cognizant of the fact that we are essentially working with thresholds. This means that structural change can be obscured depending on the location of an airport in a given cluster. To elaborate, take for instance an airport that lies on the border between cluster C5 and C6 which suffered a loss of connectivity and now moved to the border between cluster C5 and C4. Depending on the size of the intra-cluster variation, this change can be seen as a significant change in connectivity, but it will not be picked up by our model, as the airport still resides in cluster C5.

Another problem related to the clustering was made apparent in the finding of an optimal cluster center for all the weeks in the year. Going with six clusters means that we had to make a concession where we are overfitting the data in some weeks, and underfitting the data in other weeks. Hence, depending on the research question (i.e., studying the evolution of the connectivity of an airport or region over a given period), one may best opt to not use clustering and compare the heat diffusion directly.

REFERENCES

[1] EUROCONTROL (2022). Performance Review Report of the European Air Traffic Management System in 2021

[2] Bao, X., Ji, P., Lin, W., Perc, M., & Kurths, J. (2021). “The impact of COVID-19 on the worldwide air transportation network”, *Royal Society open science*, 8(11), 210682.

[3] Li, S., Zhou, Y., Kundu, T., & Sheu, J. B. (2021). “Spatiotemporal variation of the worldwide air transportation network induced by COVID-19 pandemic in 2020”, *Transport Policy*, 111, 168-184.

[4] Sun, X., Wandelt, S., Fricke, H., & Rosenow, J. (2021). “The Impact of COVID-19 on Air Transportation Network in the United States, Europe, and China”, Sustainability, 13(17), 9656.

[5] Taylor, D., Myers, S. A., Clauset, A., Porter, M. A., & Mucha, P. J. (2017). “Eigenvector-based centrality measures for temporal networks”, Multiscale Modeling & Simulation, 15(1), 537-574.

[6] Newman, M. E. (2005). “A measure of betweenness centrality based on random walks”, Social networks, 27(1), 39-54..

[7] Berkhin, P. (2005). “A survey on PageRank computing”, Internet mathematics, 2(1), 73-120

[8] Donnat, C., Zitnik, M., Hallac, D., & Leskovec, J. (2018). “Learning structural node embeddings via diffusion wavelets”, In Proceedings of the 24th ACM SIGKDD international conference on knowledge discovery & data mining (pp. 1320-1329).

[9] Sun, X., Wandelt, S., & Linke, F. (2014). “Topological properties of the air navigation route system using complex network theory”, *Proceedings of the ICRAT*.

[10] SNAP (2018). *Spectral Wavelets for learning structural signatures in complex networks*, GitHub repository, https://github.com/snap-stanford/graphwave

[11] Caliński, T., & Harabasz, J. (1974). “A dendrite method for cluster analysis”, Communications in Statistics-theory and Methods, 3(1), 1-27.

[12] EUROCONTROL (2022) Active and inactive aircraft parked across EUROCONTROL area [Online] Available: https://ansperformance.eu/covid/acft_ground/