

Efficient On-line Computation of Visibility Graphs

Delia Fano Yela, Florian Thalmann, Vincenzo Nicosia, Dan Stowell, and Mark Sandler

Abstract—A visibility algorithm maps time series into complex networks following a simple criterion. The resulting visibility graph has recently proven to be a powerful tool for time series analysis. However its straightforward computation is time-consuming and rigid, motivating the development of more efficient algorithms. Here we present a highly efficient method to compute visibility graphs with the further benefit of flexibility: on-line computation. We propose an encoder/decoder approach, with an on-line adjustable binary search tree codec for time series as well as its corresponding decoder for visibility graphs. The empirical evidence suggests the proposed method for computation of visibility graphs offers an on-line computation solution at no additional computation time cost. The source code is available online.

Index Terms—on-line, visibility graphs, binary trees, networks.

I. INTRODUCTION

IN the last decade, several methods to map time series into graphs have been proposed under the hypothesis that, appropriate graph representations can preserve the original time series information while providing alternatives to deal with non-linearity and multi-scale issues typical of complex signals [5], [25]. This line of research represents a bridge between nonlinear signal analysis and complex network theory, and has been successfully applied to extract meaningful information from a variety of different systems in physics [27], [16], finance [8], [15], [24], engineering [28], and neuroscience [2], [3].

The most notable algorithms to construct a graph from an ordered sequence of data points are either based on correlation [21], [1], [29], recurrence [6], [4], [7], dependence [22], [18], or visibility [13]. However, the visibility algorithms proposed by Lacasa et al. [13], [19] are amongst the most popular as they provide a deterministic and non-parametric symbolisation of a time series preserving full information of its linear and non-linear correlations. Such visibility algorithms can also effectively deal with non-stationary signals and are deemed computationally efficient. In consequence, visibility graphs have found numerous applications in diverse fields including image processing [12], [11], number theory [14], finance [15], [9], and neuroscience [26].

The straightforward computation of visibility graphs presents a worst case time complexity quadratic in the length of the series. Even though such complexity should not be an issue for medium-sized series ($10^4 - 10^5$ points), it remains inefficient for longer time series. Therefore, faster algorithms have been proposed employing a ‘Divide & Conquer’ (DC) approach, reducing the average-case time complexity to $O(n \log n)$ [17].

Both of these approaches comprising the current existing methods to compute visibility graphs, are off-line algorithms, as they require all the data points in the time series to be available before the graph is constructed. Consequently, the integration of new data points normally requires to recompute the visibility graph from scratch, representing a major shortcoming limiting the real-world applications of visibility graphs.

In this paper we present, to the best of our knowledge, the first on-line algorithm to compute visibility graphs efficiently. The proposed algorithm employs an ‘encoder/decoder’ approach by means of a binary search tree representation of the time series (or any ordered sequence of data points). In particular, the time series is encoded into a binary search tree that can be updated every time a chunk of time series is available by merging its corresponding binary search trees. The resulting binary search tree can subsequently be decoded into a visibility graph when required. This introduced flexibility comes at no significant computational cost as the presented method shares the computational complexity of the current fastest visibility algorithm (DC).

II. VISIBILITY GRAPHS

A visibility graph is obtained from an ordered sequence of values by associating each datum to a node and connecting two nodes with an edge if the corresponding data points are visible from each other. A point a is visible from the point b if one can draw a straight line from a to b without passing underneath any intermediate points. In this paper we will consider visibility as a symmetric relation, so that the resulting visibility graphs are undirected.

The natural visibility criterion (NV) allows the visibility line between a and b to take any slope, whereas the horizontal visibility criterion (HV) is restricted to horizontal lines, as shown in Figure 1.f. More precisely, given a time series

$$y = f(t)$$

of length n , two points (t_a, y_a) and (t_b, y_b) are said to be naturally visible if every intermediate point (t_c, y_c) , such that $t_a < t_c < t_b$, fulfills the following simple geometrical criterion:

$$y_c < y_a + (y_b - y_a) \frac{t_c - t_a}{t_b - t_a}$$

This natural visibility criterion will therefore establish the connections between nodes in the resulting natural visibility graph (NVG).

One can analogously map a time series into a horizontal visibility graph (HVG) where two points (t_a, y_a) and (t_b, y_b) are said to be horizontally visible if :

$$y_a, y_b > y_c \quad \forall c \text{ such that } t_a < t_c < t_b$$

From the definition of visibility it immediately follows that, for a set visibility criterion, the visibility graph associated to a given time series is unique. Moreover, any two subsequent data points of the time series are always connected by an edge, thus visibility graphs are connected and Hamiltonian [20]. In addition, visibility graphs are also invariant to re-scaling on both horizontal and vertical axes (i.e., the first point on either side of a node i remains visible from i no matter how far apart they are), and invariant to vertical and horizontal translations (i.e., only the relative values of point determine visibility relations).

In Figure 1.f. we show both the natural and horizontal visibility criteria at work on an arbitrary time series. Notice that horizontal visibility is a more stringent criterion than natural visibility, meaning that if two points are horizontally visible then they are also trivially visible when using the natural visibility criterion. Consequently, the horizontal visibility

graph of a time series is always a sub-graph of the natural visibility graph associated to the same time series.

III. STATE OF THE ART

A straightforward approach to compute visibility graphs consists in checking whether any of the points of the time series is visible or not from every other point. This corresponds to evaluating the visibility criteria for every pair of points in the time series. Since we consider visibility as a symmetric relation, the total number of checks needed to obtain a visibility graph of a time series of n data points is equal to $n(n-1)/2$, corresponding to a $O(n^2)$ time complexity.

In the case of horizontal visibility, one can take a step further and safely assume that no point after a value larger than the current value t_a will be horizontally visible from t_a . This observation effectively reduces the time complexity of the construction to $O(n \log(n))$ and, in the case of noisy (stochastic or chaotic) signals, it can be proved that this algorithm has an average-case time complexity $O(n)$ [20]. Nevertheless, all pairs of points need to be checked in the case of natural visibility. From now on, this simple approach will be referred to as the basic method for both natural and horizontal visibility computation ¹.

As an improved alternative for visibility computation, Lan et al. presented a ‘Divide & Conquer’ approach [17]. This algorithm reduces the average case time complexity of the construction of the natural visibility graph to $O(n \log(n))$ and it significantly reduces computation time for most balanced time series.

The basic idea behind the ‘Divide & Conquer’ algorithm is related to the horizontal visibility optimisation mentioned above. Once the maximum value M of the time series is known, one can safely assume that the points on the right of M will not be naturally visible from the points on the left of M (the point M is effectively acting as a wall between the two sides of the time series). The same argument is then applied recursively on the two halves of the time series separated by M , where the local maxima subsequently found at each level are connected with an edge to the maxima at the level immediately above them. From now on, this improved method will be referred to as ‘Divide & Conquer’ (or DC for short).

Both the basic method and DC are off-line approaches, meaning that they require all the points of the time series to be accessible at the beginning of the computation. This rigid requirement limits the applicability of visibility graphs, specially in fields like telecommunications or finance, where there is a constant incoming flow of new data to be processed and assimilated. Moreover, in such big data scenarios, one tends to favour an initial overall high level analysis that will reveal the need for further processing. This work-flow would benefit from dynamic algorithms unlike the ones presented above.

¹The original Fortran 90 implementations of basic algorithms to construct visibility graphs can be found at <http://www.maths.qmul.ac.uk/~lacasa/Software.html>

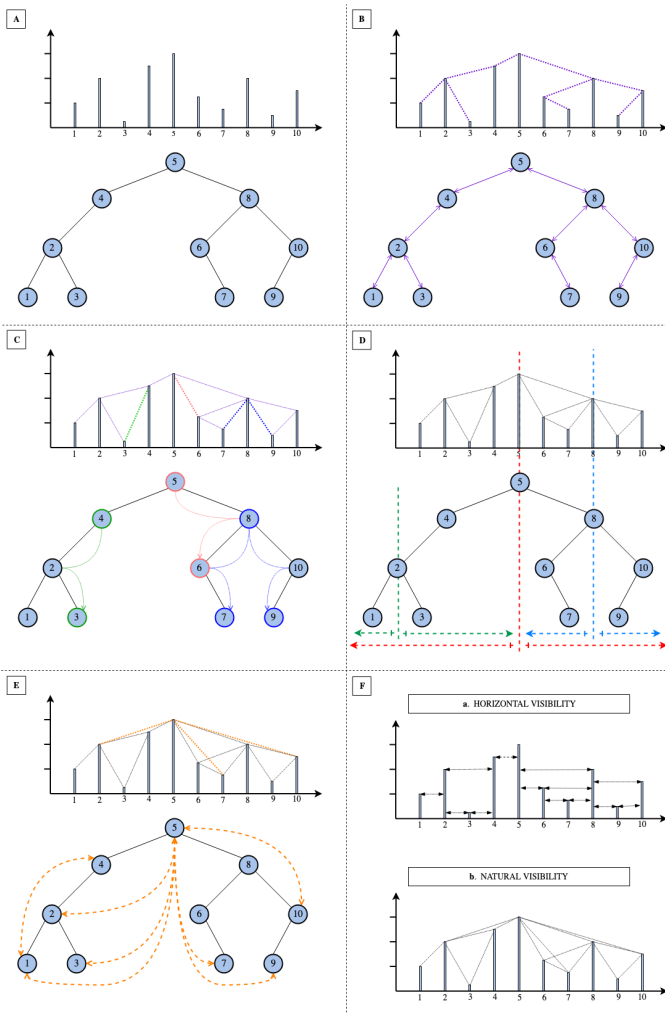


Fig. 1. Representation of the different steps of the proposed algorithm for visibility graphs computation. In section A, the sample time series and its correspondent maximum binary search tree. Section B represents the connections deduced by the first connectivity rule. The second and third connectivity rules are illustrated in section C and D respectively. Section E shows the remaining checks needed to ascertain natural visibility. Finally, section F reports the horizontal and natural visibility graph associated to the original time series.

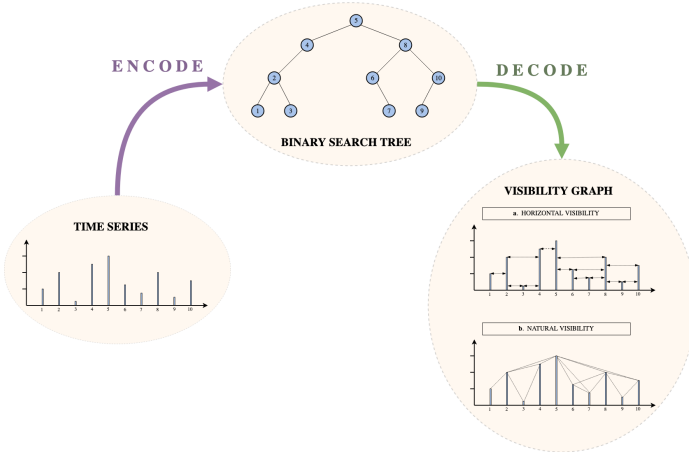


Fig. 2. Illustration of the encode/decode approach of the proposed method to calculate visibility graphs.

IV. PROPOSED METHOD: BINARY SEARCH TREE CODEC

Here we propose a new method to compute visibility graphs on-line based on an encoding/decoding approach. In our method, the necessary visibility information is first encoded into an appropriately constructed binary search tree, and then successively decoded into a visibility graph when needed, as shown in Figure 2.

A. Encoding - Maximum Binary Search Tree

The construction of a maximum binary search tree is fairly straightforward and its corresponding pseudo-code is shown in Algorithm 1. The first step is to sort the given time series in descending order of values, while storing the original position of each value in the time series. From now on, we will refer to the original positions as indices (i.e. t) and to the values of the time series simply as values (i.e. $y(t)$). In the case of repeated values in the sequence, the first encountered index will come first while sorting.

Once we have a list of values sorted in descending order, together with the corresponding indices, we follow the standard procedure to build a binary search tree based on the indices. Every entry in the index list will be a node and each node has a left and right child, as shown in the data structure proposed in Algorithm 1 (i.e., *Node*). The first node of the binary tree (the one with no *parent*) is called *root*. In our case, the root will be the index of the datum corresponding to the maximum value in the time series, which is also the first entry in the index list.

The next index, corresponding to the point with the largest value smaller than the maximum, will then be added to the tree. If its index is smaller than the root, it will become the left child of root, while if its index is larger than the root it will become the right child of root (see function *add* in Algorithm 1). The next index to add will start off being compared to the root; if its smaller, it will travel to the left of the tree and, if its bigger, to the right. It will continue descending the tree in this manner until it finds an empty spot. We continue adding the indices in the list accordingly

(see function *build_tree* in Algorithm 1) until there are no more data points (i.e. indices) to add.

In the case of the sample time series in Figure 1.a, the maximum is in position 5 and will therefore be the *root* of the binary tree. The point whose value is immediately smaller than the maximum is in position 4 (less than 5), so it will become the left child of the root. The third point in the list is in position 2, and will travel down the tree on the left-most branch (as it is smaller than both 5 and 4). The right branch of the tree is populated by the fourth point (in position 8), whose index is bigger than the *root*. In Figure 1.A one may appreciate the correlation between the time series and its associated binary tree structure. The visibility information captured by such tree may also now be more apparent.

The time complexity of the procedure needed to encode the time series into the maximum binary search tree is $O(S + T)$ where $O(S)$ is the time complexity of sorting the series and $O(T)$ is the time complexity of the algorithm to construct the binary search tree. Sorting by comparisons is known to be $O(n \log n)$ (e.g., by using either MergeSort or QuickSort), while constructing a binary search tree costs on average $O(n \log n)$. Hence the overall average-case time complexity of the encoding step is $O(n \log n)$.

```

Node {
  index : float # x, input, argument
  value : float # f(x), output
  left  : Node  # left child subtree
  right : Node  # right child subtree
}

def buildTree(values : {float}, indexes: {float}):

  root ← Node()

  sorted_values = sort_descending(values)
  sorted_indexes = indexes[getIndex(sorted_values)]

  for (i, v) in (sorted_indexes, sorted_values):
    root.add(Node(index = i, value = v))

  return root

def add(self : {Node}, node : {Node}):

  if self is empty :
    self.index = node.index
    self.value = node.value
  else:
    if node.index < self.index:
      self.left.add(node)
    else:
      self.right.add(node)

```

Algorithm 1. Pseudocode of the algorithm used to build a maximum binary search tree

B. Decoding - Connectivity Rules

The structure of the maximum binary search tree encodes sufficient information about the time series to allow to efficiently construct the corresponding horizontal visibility graph.

The decoding procedure is based on the following connectivity rules, also illustrated in Figure 1 :

- 1) All the nodes connected by an edge in the maximum binary search tree are visible to each other and therefore connected in the visibility graph (Figure 1.B);
- 2) Each node of the maximum binary search tree sees all the nodes in the left-most branch of the sub-tree rooted at its right child, as well as all the nodes in the right-most branch of the sub-tree rooted at its left child (Figure 1.C);
- 3) The nodes of the left sub-tree of a node i are not visible from the nodes of the right sub-tree of node i (Figure 1.D)

Note that, if there are no adjacent repeating amplitudes, the horizontal visibility graph is fully determined by these connectivity rules. In particular, when checking the connectivity rules, we simply skip a node if it has the same value as the current node. One can think of adjacent points with equal value as an interconnected ‘super node’, which takes the smallest index value when ‘looked’ from the left and the biggest index value when ‘looked’ from the right or from above.

Since the horizontal visibility decoding will always be fully determined by the three connectivity rules above, its time complexity is the sum of the time complexity of the rules. Essentially, each rule can be reduced to a series of look-ups in a binary search tree, and each look-up operation has time complexity $O(\log(n))$ in a balanced tree. These connectivity rules are applied to every node in the tree, and so the overall time complexity of decoding a horizontal visibility graph is $O(n \log(n))$. This represents a major improvement over the state-of-the-art algorithms, which can ramp up to $O(n^2)$ in the worst case scenario.

The construction of the natural visibility graph, instead, requires the creation of some connections that are not captured by the three connectivity rules above. Hence, in this case we need to perform additional visibility checks (Figure 1.E). In particular, for each node i we must check the natural visibility criterion with each node in the sub-tree rooted at the right child of i and with each node in the sub-tree rooted at the left child of i . These additional checks do not modify the average-case time complexity (which remains $O(n \log n)$), but the worst-case scenario still depends on the actual structure of the time series, and yields a time worst-case time complexity $O(n^2)$ for monotonically increasing or decreasing time series.

C. Time Complexity

In order to determine the time complexity of the proposed method, we will follow the standard procedure by considering the worst-case and average-case scenarios. In both scenarios, the time complexity of the encoding stage is determined by the time complexity of the sorting algorithm used, which in general is $O(n \log(n))$, and of the construction of the binary search tree, which is $O(n \log n)$. So in both cases encoding into a binary search tree costs $O(n \log n)$.

Decoding into a horizontal visibility graph is made through the three rules explained in Figure 1B-D, which require only a visit of the binary search tree (with time complexity $O(n)$).

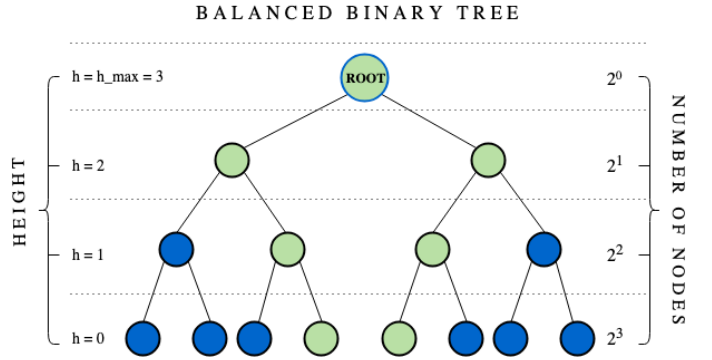


Fig. 3. Representation of a perfectly balanced tree of height 4. The nodes in green are visible to the root and this visibility can be deduced by the proposed decoder (i.e. the connectivity rules). The number of nodes at each height in a balanced tree can always be expressed in base 2.

Hence, the overall time complexity of encoding and decoding into a horizontal visibility graph is $O(n \log(n))$.

The worst case for decoding into a natural visibility graph is that of monotonically increasing, monotonically decreasing, or constant series, whose corresponding binary search trees degenerate into a line. In this case, the second and third connectivity rules are trivial, leaving only the first rule and the additional natural visibility checks. More precisely, if the tree is a line we need to check the natural visibility among $(n-1)(n-2)/2$ pairs of nodes, while the visibility of the remaining $(n-1)$ pairs of nodes is determined by the first connectivity rule. Even though this requires $(n-1)$ checks less than the basic implementation (which requires $n(n-1)/2$), the time complexity will still be $O(n^2)$ for the worst case scenario.

For the average case we assume the maximum binary search tree to be balanced. This means that the connectivity rules of the decoder will significantly reduce the overall number of visibility checks. If we consider a perfectly balanced binary tree as shown in Figure 3, the inner left branch of the right sub-tree and the inner right branch of the left sub-tree of a node are visible to the parent node. These are represented in green in Figure 3 where the root is the parent node. This means that the visibility between the root and all the rest of nodes (the ones in blue) is unknown and needs to be checked.

Therefore we can deduce that the number of remaining visibility checks for the root in a balanced tree of height h_{root} is equal to $2^{h_{\text{root}}+1} - 1 - 2h_{\text{root}}$, where $2^{h_{\text{root}}+1} - 1$ is the total number of nodes below the root while $2h_{\text{root}}$ is the number of nodes whose visibility can be deduced by the three decoding rules (green nodes). Notice that the height of the root h_{root} corresponds to the maximum height of the balanced tree h_{max} . The same reasoning applies to all the other nodes. More precisely, for a node at height h , there will be $(2^{h+1} - 1 - 2h)$ remaining visibility checks to be performed.

In order to calculate the total number of remaining visibility checks, one needs to multiply the individual expression above by the number of nodes at that height $2^{h_{\text{max}}-h}$ and sum across all heights where the checks are needed (all except the last two). Therefore, one can express the total number of remaining natural visibility checks in a perfectly balanced binary tree as

follows:

$$\sum_{h=2}^{h_{\max}} 2^{h_{\max}-h} [2^{h+1} - (2h+1)] = 2^{h_{\max}} \left[2(h_{\max}-1) - \sum_{h=2}^{h_{\max}} h2^{1-h} - \sum_{h=2}^{h_{\max}} 2^{-h} \right]$$

Since the maximum height of a balanced tree with n nodes is $h_{\max} = \log_2(n)$, the total number of operation is dominated by the first term of the expression above,

$$2^{h_{\max}} 2(h_{\max}-1) = 2n(\log_2(n)-1)$$

while the remaining terms will only introduce logarithmic corrections. In conclusion, the time complexity of the decoding for natural visibility graphs is on average $O(n \log(n))$.

The proposed method has the same average-case time complexity than the DC algorithm, thus improving on the original basic algorithm for both horizontal and natural visibility graphs. In the Experiment section below we will see that in practice our algorithm out-competes the basic algorithm and performs as well as the DC approach, with the additional property of allowing for on-line assimilation of new data points.

```
def merge(input:{Node}):

    if input is empty: return null

    r ← min_index(maxima_value(input))
    pool ← input \ {r}

    pool.append(r.left, r.right)

    for n in input \ {r} :

        for c in [n.left, n.right] :

            if sign(n.index - r.index)
                ≠ sign(c.index - r.index):

                pool.append(c)
                n.remove(c)

    return Node(
        index = r.index,
        value = r.value,
        left =
            merge({p | p ∈ pool, p.index < r.index}),
        right =
            merge({p | p ∈ pool, p.index > r.index}))
```

Algorithm 2. Pseudocode of the proposed algorithm to merge two binary trees defined by their root (class Node). The input is a list of roots to be merged.

V. ON-LINE VISIBILITY GRAPHS: MERGING BINARY TREES

Every time a node is added to an existing binary search tree, it essentially ‘travels’ down the tree, going left if smaller and right if larger, until it finds an empty space (see pseudocode function *add* in Algorithm 1). Therefore when a node is added to an existing binary tree there is no need to recalculate the tree structure from scratch. Due to the fact that the proposed

encoder is a binary search tree, there is a possibility to efficiently update it on-line.

Given a time series and its correspondent binary search tree, we would like to integrate new data points in the tree structure without recomputing it from scratch. One could process the points of the newly available batch of data individually and include them in the existing tree structure by comparing both values and indices. However, other than being a time consuming approach for large numbers of points, processing points individually fails to include useful information of both the batch and the current tree structure. As an example, in Figure 4.A, all the nodes in the batch to be added (red nodes) have larger indices than the nodes in the current tree structure (blue nodes), and so larger indices than the current root. This means, all the nodes in the batch will populate the right side of the resulting tree. If the nodes are treated individually, this information will be overlooked producing an inefficient algorithm.

Therefore, we propose to take a different approach by treating the new batch of points as an entity. More precisely, we propose to compute the binary search tree of the new nodes and *merge* it with the previous tree structure as illustrated in Figure 4. In this way, if all the new nodes indices are larger than the current root, one can include such information and produce an optimised algorithm, where potentially only one comparison is needed to merge the current with the batch tree. This is the case for real-time incoming data, as the batch’s nodes always have larger time values (indices) than the previous points in the time series.

Furthermore, the proposed merge approach covers both append and insert operations, illustrated in Figure 4.A and Figure 4.B respectively. In terms of time series representation, this means one could update the binary tree codec with observations that happened later in time or with a higher time resolution. This novel introduced flexibility for visibility computation, opens the door to new applications such as big data or audio applications where the sampling rate may vary at different analysis stages.

In order to merge two trees, we propose to compare them by levels, increasing depth at every recursion of the *merge* function outlined in Algorithm 2. The comparison happens in two steps: firstly the node values at a level are compared to determine which node will occupy that location in the resulting tree; secondly, the node indices are compared to determine which direction the rest of the nodes will travel down in depth.

Following the construction of the proposed binary search tree, the node with larger value will be chosen and the rest of the nodes will travel left if their indices are smaller than the chosen one and right otherwise. The nodes to be compared are the children of the chosen node with the nodes from the previous level that were not chosen; starting of by comparing the two roots of the trees to be merged. The *merge* algorithm is illustrated step by step at the top and bottom of Figure 4.

In the example in Figure 4.A, the blue and red tree are to be merged. Initially, the blue root is compared to the red root. Since the blue root has a larger value than the red root, it will be chosen to take that position in the resulting tree (i.e. the root of the resulting tree). The red root will then travel down

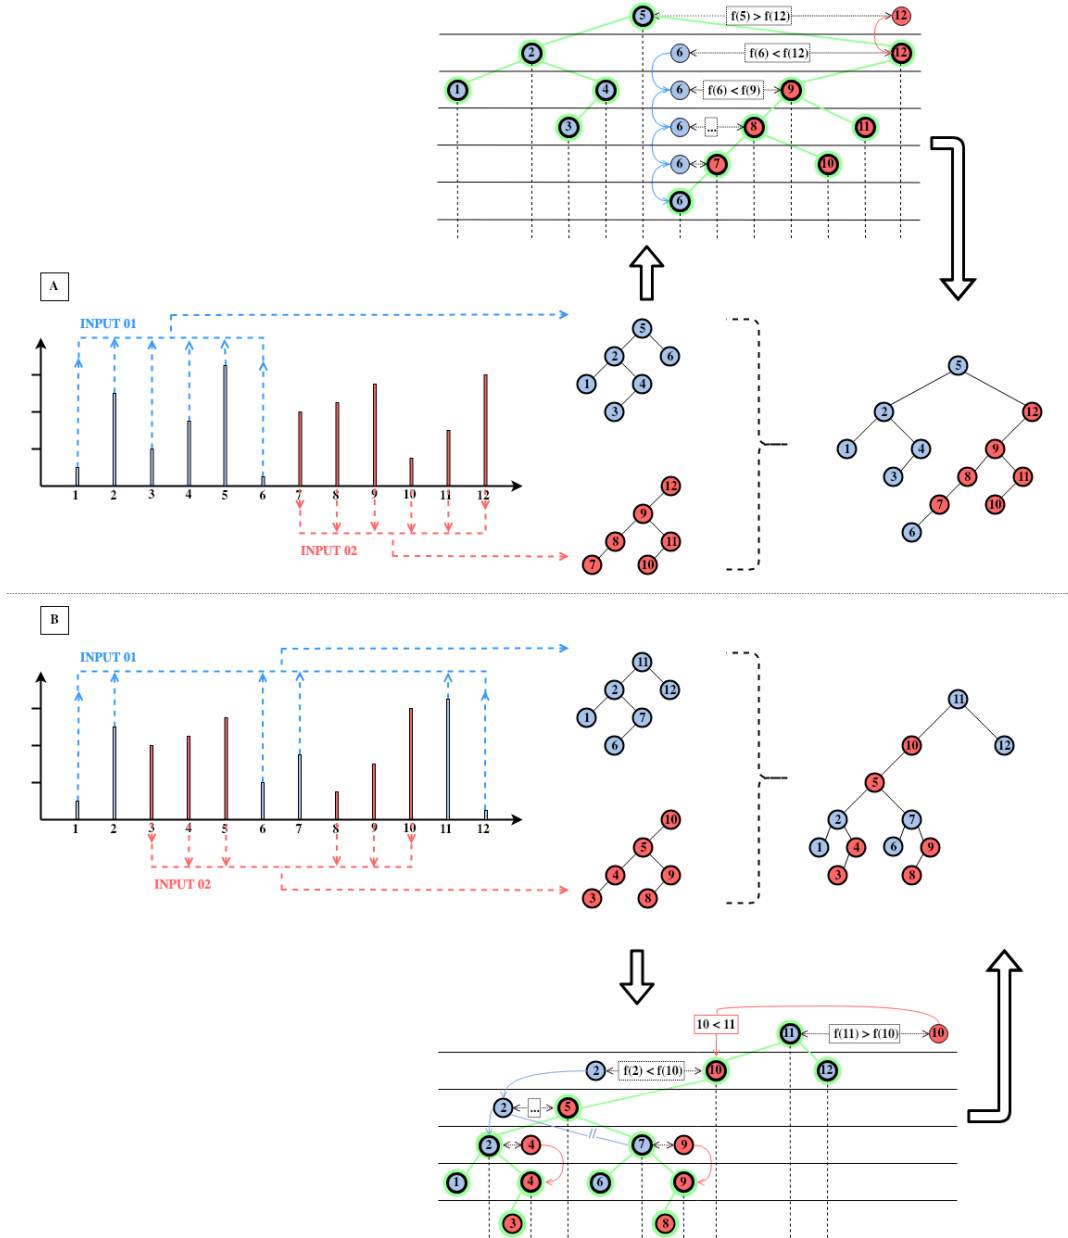


Fig. 4. Visual representation of the proposed method to merge two maxima binary trees, covering both append (A) and insert (B) operations. This corresponds to an on-line scenario where a new batch (red) needs to be incorporated to an existing structure (blue).

the right branch as its index is larger than the chosen blue root, leaving the left branch of the chosen blue root untouched.

Consequently the right blue child is to be compared with the red root. In this case, the red root has a larger value and so it will take the right branch position in the resulting tree. Now it is the turn for the right blue child to descend down the red tree. Since the blue child happens to be the lowest value in the series, it will just descend layers following the binary search tree rules until it reaches an empty spot.

Usually, as one may observe in Figure 4, the children of the nodes that travel down in depth are not included in the level comparison. However, when new data is to be inserted to the existing series, the child of the node traveling down could have an index corresponding to the other branch of the resulting tree. In this case, the connection between the node

and that child will be broken thereafter. For example, in Figure 4.B., this situation takes place in layer 3, where Node 7, the child of Node 2, belongs to the right branch of Node 5, unlike its parent.

VI. NUMERICAL EXPERIMENTS

In this section we present empirical results in order to show how the proposed visibility algorithm compares to the state of the art. All the code related to this paper and necessary to run the following experiments is implemented in Python 2.7 and freely available online ². The machine used in the simulations is an early 2015 MacBook Pro Retina with a 2.9GHz Intel Core i5 processor and 16GB of RAM.

²Available at <https://github.com/delialia/bst>

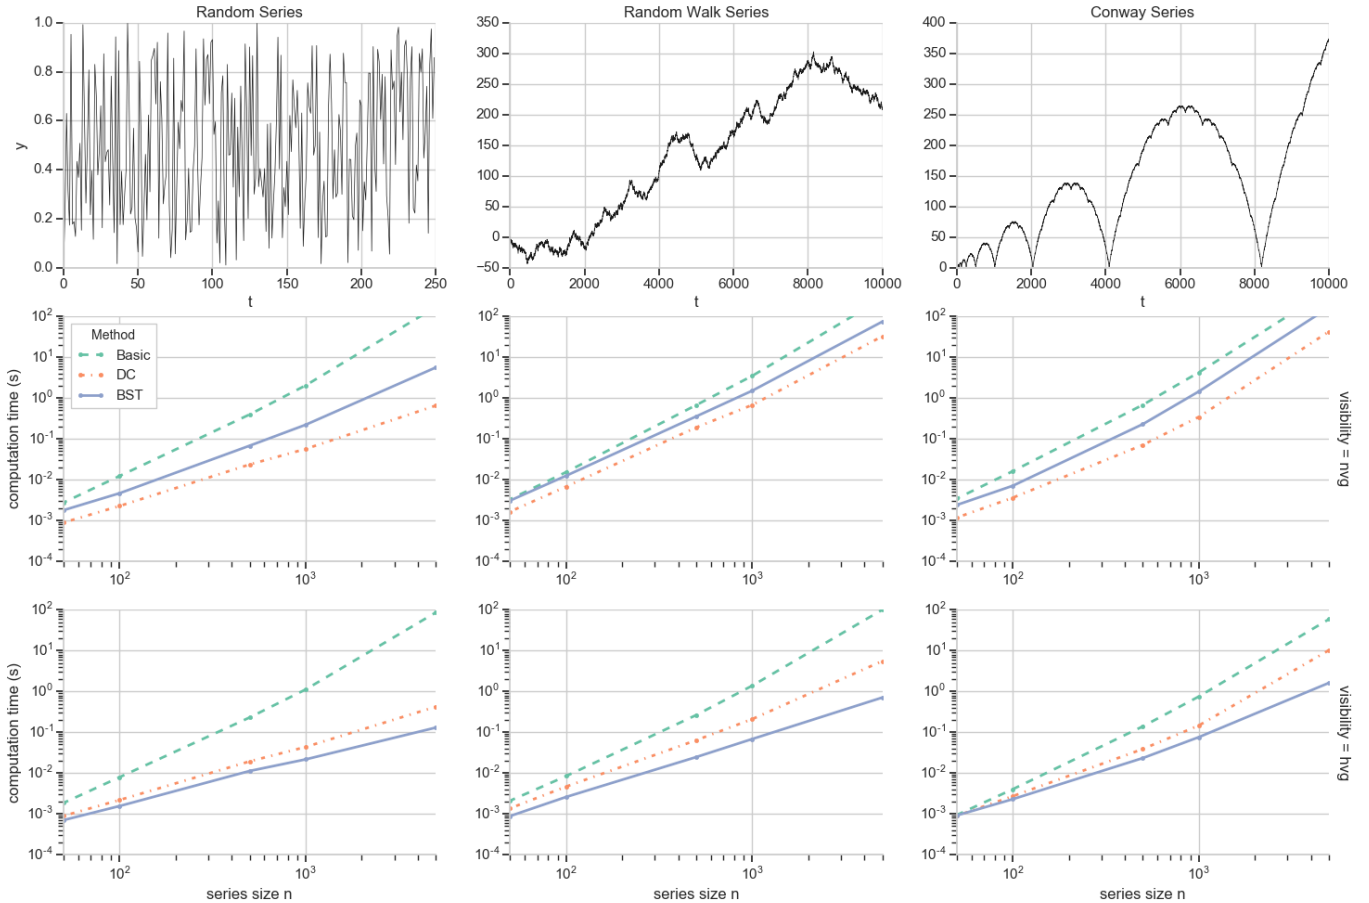


Fig. 5. Computation time of the natural visibility graph (nvg, second row) and horizontal visibility graph (hvg, third row) of different time series (examples on first row) using the current visibility algorithms: Basic, Divide & Conquer (DC), and the proposed binary search tree (BST) method. Each point at every series size is the mean of the computation time for 10 series of that size.

To put the presented algorithm into context [17], in Figure 5 we report the computation time needed by current visibility algorithms on different synthetic time series of increasing length. Since the actual efficiency of each algorithm depends to some extent on the character of the original time series, we considered uniform random noise (which has no structure and on average produces almost-balanced binary search trees), a Conway series (which has a quite rich structure and corresponds to a quite unbalanced tree), and a random walk series (which represents the more realistic scenario of a signal with both structure and noise).

In the first case we observe the largest gap in computation time between the basic algorithm and the more efficient ones as it corresponds to the aforementioned average case where both algorithms (DC and the proposed one) significantly reduce the number of operations. Such differences are more prominent in the computation of the horizontal visibility graph.

Additionally, in Figure 6 we present a similar computational time analysis over real samples of speech (English language) [10] and financial data [23]. Figure 6 is particularly interesting as it clearly shows a correlation between time computation and the time series structure (please note the different scale for time computation). Even though the time computation may differ, the DC and proposed method distribution seem

to vary very little between data types in comparison to the relatively high spread observed for the basic algorithm.

The horizontal visibility computation remains stable in both the DC and proposed method, and could potentially be considered independent of the data type given a time computation scaling factor. This behaviour was expected as the proposed method is fully defined by the aforementioned connectivity rules and has average-case time complexity $O(n \log n)$.

On the other hand, Figure 6 suggests that the efficiency of the computation of natural visibility graphs is subject to wider fluctuations. The position of the maximum in the time series affects the efficiency of both the DC and the proposed method, as it will determine the number of additional visibility checks needed to obtain the natural visibility graph.

An English speech time series will typically have its maximum somewhere towards the middle section of the signal (since we rarely tend to raise our voice at the end of our speech). Therefore the speech time series proposed codec will most probably produce an almost balanced binary search tree, yielding a time complexity of $O(n \log n)$. For this reason, one may observe a wider gap in computation time between the basic method and the faster alternatives for the speech data in Figure 6 than for the financial time series.

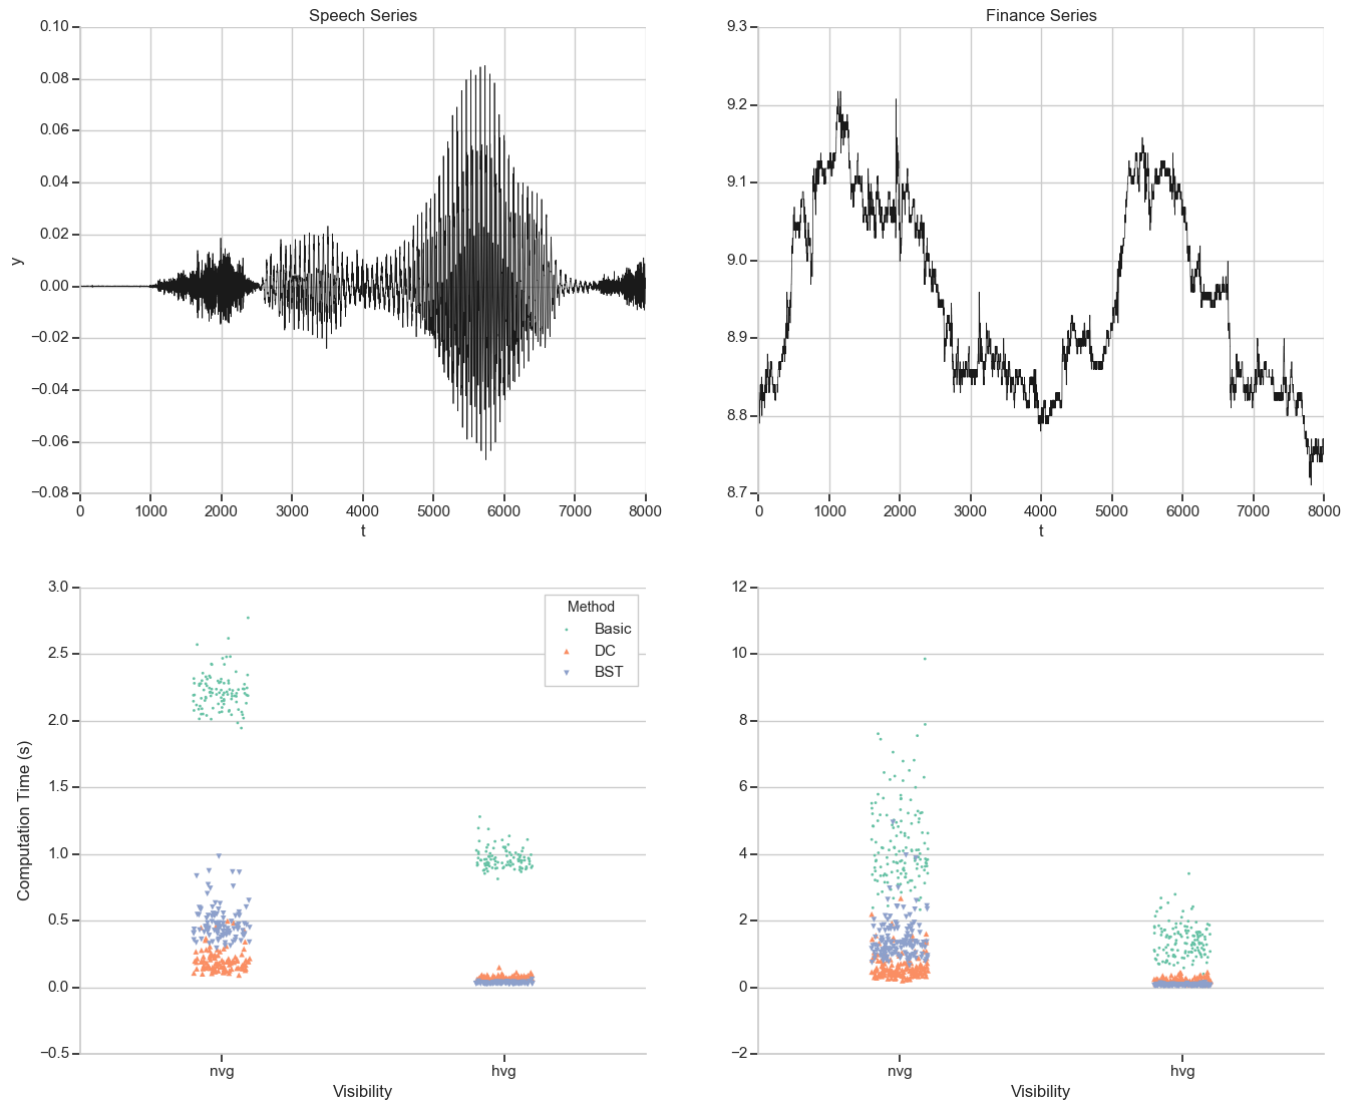


Fig. 6. Current and proposed visibility algorithms computation time for 100 speech and finance time series of 1000 points. The speech time series are sampled from the training TIMIT dataset [10]. The finance time series corresponds to the 2013 quarterly data used in [23].

In terms of computation time, the proposed method and the DC one are closely related. They are both quicker than the basic implementation in both natural and horizontal visibility and they both present similar trends for increasing time series size (Figure 5). However, the proposed algorithm has proven to consistently be the quickest option for horizontal visibility graph computation. On the other hand, the DC algorithm in general does perform better than the proposed method for natural does visibility computation. Even though at this point both DC and the proposed method seem equally good of an option for fast visibility computation, the presented algorithm has the additional property of allowing on-line assimilation of new data, which is something not easily achievable in either the basic approach or the DC algorithm.

The most straightforward way to assess the on-line functionality of the proposed method is to compare it with the equivalent off-line approach. In our case, it directly relates to the binary tree codec. Given a batch of new points to be added to the time series visibility analysis, in the off-line approach,

the new batch is simply added to the time series itself and then the binary tree codec must be re-computed from scratch. In the proposed on-line approach, the next batch is encoded into its own binary tree that is then merged to the existing codec using the procedure detailed in Algorithm 2. Note that the decoding step remains the same for the on-line and off-line approach, and so the comparison will essentially be between computing a codec from scratch (off-line) and merging two codecs into a single binary search tree (on-line).

Figure 7 shows how much quicker the computation of the on-line method (codec for new data + merging) is in comparison to the computation time of the off-line approach (codec from scratch), for different time series and batch sizes. In particular, the on-line approach is always better if the new batch to be added is equal or bigger than the existing time series, especially for large time series.

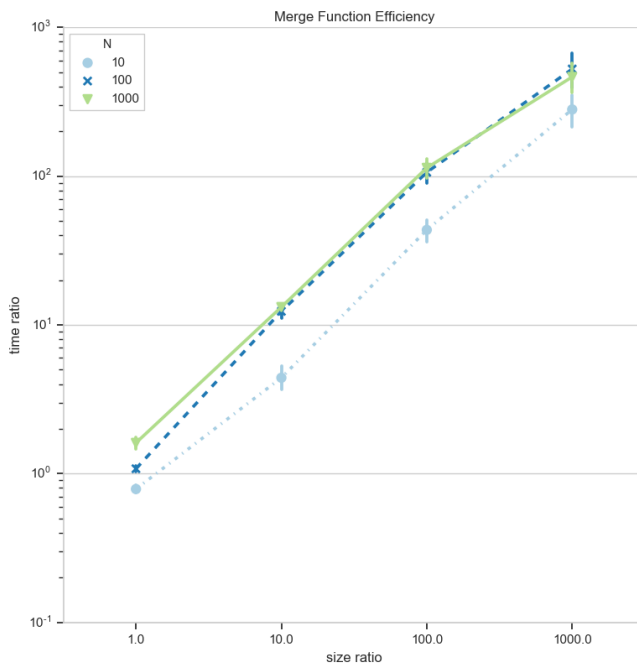


Fig. 7. Given a random time series (size L) and a batch of new random points (size N) to be added to it, this plot shows the advantage, in terms of computation time, of the proposed on-line approach versus the off-line alternative. The proposed method computation time is the time it takes to build the maxima tree of the new points and merge it with the existing time series tree (i.e. $t_{on-line}$). The off-line alternative computation time is the time it takes to build a new maxima binary tree from scratch including the new points to the time series (i.e. $t_{off-line}$). The time ratio is the \log scale of $t_{off-line}/t_{on-line}$, how much quicker the proposed method is. The size ratio is L/N , how much bigger the time series is compared to the batch to be added. Both append and insert scenarios are represented here, 10 random cases of each were computed. The point in the graph is the mean of these 20 cases and its uncertainty is captured by the error bars.

VII. CONCLUSION

The proposed visibility algorithm based on an encoder/decoder approach is, up to the authors' knowledge, the first efficient on-line algorithm to compute visibility graphs. The analysis and the numerical experiments shown in the paper confirm that the proposed algorithm represents a substantial improvement over the state-of-the-art for horizontal visibility computation, and is on par with the most efficient natural visibility algorithm (i.e. DC) available. Moreover, the procedure to assimilate new data by means of merging the corresponding binary search tree encoding into the existing tree allows for efficient on-line computation of visibility graphs, and represents a substantial speed-up with respect to the existing off-line algorithms. This novel on-line capability broadens the applications for visibility graphs at no additional computational cost.

REFERENCES

- [1] G. Bonanno, F. Lillo, and R. Mantegna. High-frequency cross-correlation in a set of stocks. *Quantitative Finance*, 1(1):96–104, Jan 2001.
- [2] E. Bullmore and O. Sporns. Complex brain networks: graph theoretical analysis of structural and functional systems. *Nat Rev Neurosci*, 10(3):186–198, Mar. 2009.
- [3] E. Bullmore and O. Sporns. The economy of brain network organization. *Nat Rev Neurosci*, Apr 2012.

- [4] J. F. Donges, R. V. Donner, and J. Kurths. Testing time series irreversibility using complex network methods. *EPL (Europhysics Letters)*, 102(1):10004, apr 2013.
- [5] R. V. Donner, M. M. Small, J. F. Donges, N. Marwan, Y. Zou, R. Xiang, and J. Kurths. Recurrence-based time series analysis by means of complex network methods. *International Journal of Bifurcation and Chaos*, 21(04):1019–1046, 2011.
- [6] R. V. Donner, Y. Zou, J. F. Donges, N. Marwan, and J. Kurths. Recurrence networks: a novel paradigm for nonlinear time series analysis. *New Journal of Physics*, 12(3):033025, 2010.
- [7] J. H. Feldhoff, R. V. Donner, J. F. Donges, N. Marwan, and J. Kurths. Geometric signature of complex synchronisation scenarios. *EPL (Europhysics Letters)*, 102(3):30007, may 2013.
- [8] P. Fiedor. Networks in financial markets based on the mutual information rate. *Phys. Rev. E*, 89:052801, May 2014.
- [9] R. Flanagan and L. Lacasa. Irreversibility of financial time series: A graph-theoretical approach. *Physics Letters A*, 380(20):1689 – 1697, 2016.
- [10] J. S. Garofolo, L. F. Lamel, W. M. Fisher, J. G. Fiscus, and D. S. Pallett. Darpa timit acoustic-phonetic continuous speech corpus cd-rom. nist speech disc 1-1.1. *NASA STI/Recon technical report n*, 93, 1993.
- [11] J. Iacovacci and L. Lacasa. Visibility graphs for image processing. *IEEE Transactions in Pattern Analysis and Machine Intelligence*, in press.
- [12] L. Lacasa and J. Iacovacci. Visibility graphs of random scalar fields and spatial data. *Phys. Rev. E*, 96:012318, Jul 2017.
- [13] L. Lacasa, B. Luque, F. Ballesteros, J. Luque, and J. C. Nuno. From time series to complex networks: The visibility graph. *Proceedings of the National Academy of Sciences*, 105(13):4972–4975, 2008.
- [14] L. Lacasa, B. Luque, I. Gmez, and O. Miramontes. On a dynamical approach to some prime number sequences. *Entropy*, 20(2), 2018.
- [15] L. Lacasa, V. Nicosia, and V. Latora. Network structure of multivariate time series. *Scientific reports*, 5:15508, 2015.
- [16] L. Lacasa and R. Toral. Description of stochastic and chaotic series using visibility graphs. *Phys. Rev. E*, 82:036120, Sep 2010.
- [17] X. Lan, H. Mo, S. Chen, Q. Liu, and Y. Deng. Fast transformation from time series to visibility graphs. *Chaos: An Interdisciplinary Journal of Nonlinear Science*, 25(8):083105, 2015.
- [18] W. Liao, J. Ding, D. Marinazzo, Q. Xu, Z. Wang, C. Yuan, Z. Zhang, G. Lu, and H. Chen. Small-world directed networks in the human brain: Multivariate granger causality analysis of resting-state fmri. *NeuroImage*, 54(4):2683 – 2694, 2011.
- [19] B. Luque, L. Lacasa, F. Ballesteros, and J. Luque. Horizontal visibility graphs: Exact results for random time series. *Physical Review E*, 80(4):046103, 2009.
- [20] B. Luque, L. Lacasa, F. Ballesteros, and J. Luque. Horizontal visibility graphs: Exact results for random time series. *Phys. Rev. E*, 80:046103, Oct 2009.
- [21] R. Mantegna. Hierarchical structure in financial markets. *Eur. Phys. J. B*, 11(1):193?197, Sep 1999.
- [22] D. Marinazzo, M. Pellicoro, and S. Stramaglia. Kernel method for nonlinear granger causality. *Phys. Rev. Lett.*, 100:144103, Apr 2008.
- [23] N. Musmeci, V. Nicosia, T. Aste, T. Di Matteo, and V. Latora. The multiplex dependency structure of financial markets. *Complexity*, 2017, 2017.
- [24] N. Musumeci, V. Nicosia, T. Aste, T. D. Matteo, and V. Latora. The multiplex dependency structure of financial markets. *Complexity*, 2017(9586064), 2017.
- [25] A. M. Nuñez, L. Lacasa, J. P. Gomez, and B. Luque. Visibility algorithms: A short review. In *New Frontiers in Graph Theory*. InTech, 2012.
- [26] S. Sannino, S. Stramaglia, L. Lacasa, and D. Marinazzo. Visibility graphs for fmri data: Multiplex temporal graphs and their modulations across resting-state networks. *Network Neuroscence*, 1(3):208–221, 2017.
- [27] A. H. Shirazi, G. R. Jafari, J. Davoudi, J. Peinke, M. R. R. Tabar, and M. Sahimi. Mapping stochastic processes onto complex networks. *Journal of Statistical Mechanics: Theory and Experiment*, 2009(07):P07046, 2009.
- [28] I. G. Torre, B. Luque, L. Lacasa, J. Luque, and A. Hernández-Fernández. Emergence of linguistic laws in human voice. *Scientific reports*, 7:43862, 2017.
- [29] Y. Yang and H. Yang. Complex network-based time series analysis. *Physica A: Statistical Mechanics and its Applications*, 387(5):1381 – 1386, 2008.