



PONTIFICIA UNIVERSIDAD CATÓLICA DE CHILE  
ESCUELA DE INGENIERÍA

# **TIME SERIES VARIABILITY TREE FOR FAST LIGHT CURVE RETRIEVAL**

**LUCAS VALENZUELA PUGH**

Thesis submitted to the Office of Research and Graduate Studies  
in partial fulfillment of the requirements for the degree of  
Master of Science in Engineering

Advisor:  
KARIM PICHARA

Santiago de Chile, September 2017

© MMXVII, LUCAS VALENZUELA PUGH



PONTIFICIA UNIVERSIDAD CATÓLICA DE CHILE  
ESCUELA DE INGENIERÍA

# **TIME SERIES VARIABILITY TREE FOR FAST LIGHT CURVE RETRIEVAL**

**LUCAS VALENZUELA PUGH**

Members of the Committee:

KARIM PICHARA

DENIS PARRA

MARCELO MENDOZA

FRANCISCO FÖRSTER

JOSÉ MIGUEL AGUILERA

Thesis submitted to the Office of Research and Graduate Studies  
in partial fulfillment of the requirements for the degree of  
Master of Science in Engineering

Santiago de Chile, September 2017

© MMXVII, LUCAS VALENZUELA PUGH

*Gratefully to my mother, my siblings  
and Ivania*

## TABLE OF CONTENTS

LIST OF FIGURES	vi
LIST OF TABLES	vii
ABSTRACT	viii
RESUMEN	ix
1. INTRODUCTION	1
2. BACKGROUND THEORY AND RELATED WORK	5
2.1. Clustering and $k$ -medoids . . . . .	5
2.2. Time Warp Edit Distance . . . . .	6
2.3. Trees . . . . .	7
2.4. Text retrieval review . . . . .	8
2.5. From the Vocabulary Tree to the Variability Tree . . . . .	11
2.6. Related Work . . . . .	11
3. DESCRIPTION OF THE METHOD	13
3.1. Building the <i>Variability Tree</i> . . . . .	15
3.1.1. Standardise light curves . . . . .	15
3.1.2. Sample subsequences . . . . .	15
3.1.3. Build the tree structure . . . . .	16
3.1.4. Populate the tree . . . . .	17
3.1.5. Build matrix representation . . . . .	19
3.2. Data retrieval: querying the tree . . . . .	21
3.2.1. Standardise query light curve and extract subsequences . . . . .	22
3.2.2. Add query subsequences to tree . . . . .	22
3.2.3. Build query vector representation . . . . .	22
3.2.4. Reduce matrix and calculate similarity . . . . .	23

3.2.5. Calculate distances and sort ranking . . . . .	24
4. DATA . . . . .	25
4.1. Datasets . . . . .	25
4.1.1. Ogle III Catalog of Variable Stars (OGLE) . . . . .	25
4.1.2. The MACHO project . . . . .	26
4.1.3. The Kepler Mission (KEPLER) . . . . .	27
4.2. Query-sets . . . . .	27
5. IMPLEMENTATION . . . . .	29
5.1. Software stack . . . . .	29
5.2. Parameters . . . . .	29
6. EXPERIMENTAL RESULTS . . . . .	31
6.1. Experimental set-up No. 1: measuring ranking quality . . . . .	31
6.2. Experimental set-up No. 2: Large datasets . . . . .	35
7. CONCLUSIONS . . . . .	38
REFERENCES . . . . .	39

## LIST OF FIGURES

2.1	The generic tree data structure with its main parts. . . . .	7
3.1	The <i>Variability Tree</i> building process. . . . .	14
3.2	The <i>Variability Tree</i> querying process. . . . .	14
3.3	Tree construction through top-down hierarchical $k$ -medoids. . . . .	17
3.4	Subsequences are extracted from every light curve in the data set using a sliding window. . . . .	18
3.5	Variability Tree after being populated with full dataset. . . . .	20
6.1	P@10 distribution by class for OGLE dataset. . . . .	34
6.2	P@10 distribution by class for MACHO dataset. . . . .	34
6.3	Confusion matrices for OGLE and MACHO datasets. . . . .	35
6.4	Sample queries on MACHO and OGLE data sets. . . . .	37

## LIST OF TABLES

2.1	Text retrieval concepts exemplified. . . . .	10
4.1	Class distribution of OGLE labelled set. . . . .	25
4.2	Class distribution of MACHO labelled set. . . . .	26
4.3	Class distribution of query-sets built for OGLE, MACHO and KEPLER datasets. . . . .	28
5.1	Values of parameters used at experiments. . . . .	29
6.1	Precisions at k for rankings obtained in experimental set-up No. 1. . . . .	33
6.2	Query time (in seconds) for experimental set-up No. 1. . . . .	33
6.3	Average time in seconds for queries in experiments for unlabelled datasets. . . . .	36

## ABSTRACT

During the last ten years, a considerable amount of effort has been made to develop algorithms for the automatic classification of variable stars. That has been primarily achieved by applying machine learning methods to photometric datasets where objects are represented as light curves. Machine learning classifiers require training sets to learn the underlying patterns that allow the separation among classes. Unfortunately, building training sets for light curves is an expensive process that demands a lot of human efforts. In addition, many times astronomers pursue the identification of a particular type of variable stars, where is not reasonable to wait for training sets that are formed mostly by other classes of variability.

In order to facilitate the creation of training sets and identification of particular light curves, we develop a fast light curve search engine, that allows finding the most similar light curves to a specific target in a data set, without the need of labelled data. We also present a novel similarity function specific for light curves, based on a novel data structure that allows scaling the search over the entire dataset of unlabelled objects. Experiments show that our model performs fast queries, achieves high precision and scales up to massive amounts of light curves.

**Keywords:** light curves, similarity search, information retrieval.



## RESUMEN

Durante los últimos diez años, se han hecho esfuerzos considerables por desarrollar algoritmos para la clasificación automática de estrellas variables. Esto ha sido logrado principalmente mediante la aplicación de métodos de aprendizaje de máquina en conjuntos de datos donde los objetos celestes son representados como curvas de luz. Estos clasificadores necesitan conjuntos de entrenamiento para aprender los patrones subyacentes que permiten la separación entre clases. Desgraciadamente, construir conjuntos de entrenamiento para curvas de luz es un proceso costoso que demanda cantidades considerables de esfuerzo humano. Además, muchas veces los astrónomos buscan identificar un tipo particular de estrellas variables. En este caso, no es razonable esperar la disponibilidad de conjuntos de entrenamiento formados mayoritariamente por otras clases de variabilidad.

Con el propósito de facilitar la creación de conjuntos de entrenamiento y la identificación de curvas de luz particulares, desarrollamos un motor de búsqueda eficiente, que permite encontrar en un conjunto de datos las curvas de luz más similares a una curva objetivo específica, sin la necesidad de datos etiquetados. Además presentamos una función de similaridad específica para curvas de luz, basada en una novedosa estructura de datos, que permite escalar la búsqueda a grandes conjuntos de objetos no etiquetados. Nuestros experimentos muestran que nuestro modelo ejecuta búsquedas rápidas, obtiene alta precisión y escala a grandes cantidades de datos.

**Palabras Claves:** curvas de luz, búsqueda de similaridad, recuperación de información.

## 1. INTRODUCTION

In the last decade, a considerable amount of effort has been made in automating the classification of variable stars [Bloom et al., 2012, Debosscher et al., 2007, Kim et al., 2014, 2011, Mackenzie et al., 2016, Nun et al., 2014, 2015, Pichara & Protopapas, 2013, Pichara et al., 2012, 2016, Richards et al., 2011]. This has been largely achieved by applying machine learning methods to photometric datasets where objects are represented by their light curves. Machine learning classifiers have been used for this purpose. Classifiers can automatically identify classes of unlabelled data after learning patterns from datasets where the classes—or labels—are already known. This learning stage is called *training* and the originally labelled dataset, from which the classifier learns, is consequently called *training set*.

Light curve datasets are highly heterogeneous. They present different number of observations and different lengths in time domain among exemplars. Light curves are also unevenly sampled and often present gaps, which makes difficult to apply directly general-purpose and traditional time series techniques [Huijse et al., 2014, Richards et al., 2011].

Algorithms coming from the machine learning community are not designed to work with this kind of data. Therefore, light curves have to be transformed into vector-like data before being passed to classifiers. Traditionally, to do this, a series of descriptors called *features* are extracted from the original light curves, so that each gets represented by a feature vector [Nun et al., 2015]. Features vary in nature, and there has been extensive research on light curve feature extraction [Bloom et al., 2012, Cabrera-Vives et al., 2017, Debosscher et al., 2007, Elorrieta et al., 2016, Förster et al., 2016, Huijse et al., 2014, Ivezić et al., 2008, Kim et al., 2014, 2011, Mackenzie et al., 2016, Nun et al., 2014, 2015, Pichara & Protopapas, 2013, Pichara et al., 2012, 2016, Richards et al., 2011].

Although some of these methods have proven largely successful, they rely on the existence of good quality training sets. To obtain good classification results, training data have to be as large as possible and be made up of template-quality light curves with already known

class. According to Angeloni et al. [2014], light curves have to meet the following conditions to be considered as templates: (1) a relative error of less than 1/10 of the light curve amplitude and (2) no significant gaps in the phase coverage. Also, all templates must be in the same band as the light curves in the dataset to be classified.

Angeloni et al. [2014] present the following alternatives to obtain training instances: (1) searching on previous surveys for labelled light curves that meet these conditions, (2) collecting labels by cross-matching a labelled dataset with one that contains light curves meeting the other requirements, (3) setting up an observational program to build light curves in the correspondent band, targeting previously studied stars. After obtaining labelled light curves by either of these methods, collected objects have to be carefully examined by an astronomer to determine if they meet template-quality.

The creation of training sets requires a considerable amount of effort from specialized scientists. Moreover, using these methods, the possible size of a training set is limited by the existence of stars whose class is already known and meet the above conditions. Insufficient training data decrease the possibilities of a successful application of automatic classifiers [Bhamidipaty et al., 2002].

Having a tool to explore and find light curves in big datasets would help astronomers to create much larger and better training sets. With such a tool, it would be feasible to find template-quality light curves which are similar to the ones previously obtained. If the classes of retrieved light curves are not previously known, these can be classified by an expert without having to examine large collections of data manually. Also, it would be possible to find in big datasets light curves that meet certain conditions defined by astronomers. This could be done by first generating artificial light curves that meet the desired conditions, and then, using a similarity search tool to find light curves in a real dataset which are similar to the ones previously created.

Time series similarity search is a problem that has been extensively studied [Agrawal et al., 1993, 1995, Baydogan & Runger, 2016, Chan & Fu, 1999, Keogh et al., 2001, Kontaki

& Papadopoulos, 2004, Popivanov & Miller, 2002, Rakthanmanon et al., 2013, Sart et al., 2010, Showcase et al., 1997, Wang et al., 2013, Wu et al., 2000]. However, to the best of our knowledge, most of the published work assume evenly spaced time series, which does not match the light curves context.

In this thesis, we present a novel method to perform fast similarity search on light curve catalogues. This method consists in organizing the light curve data into an efficient data structure for non-uniformly sampled time series we call *Variability Tree*. This structure is specially designed to retrieve light curves ranked by its similarity to a given target light curve. The Variability Tree allows us to create a new compressed representation of the light curves, enough to find similar variable stars just by using that representation. This method integrates both feature extraction and indexing for fast retrieval, without having to tackle both problems separately.

After the construction of the Variability Tree, to perform any light curve search given a query light curve, we just need to pass the query light curve through the Variability Tree to map it into the compressed representation and then get the most similar light curves ranked by similarity.

Although our method has been specially designed for light curves, taking into account their irregular nature, it can also be used on any time series dataset with even or unevenly sampled data and arbitrary length.

To the best of our knowledge, there are no published works that address the problem of efficient light curve similarity retrieval using raw light curve data. Today, the natural approach to light curve similarity retrieval is to perform feature extraction over the light curves, transforming the time series data into feature-vectors. These vectors can be easily indexed using data structures that allow fast nearest-neighbour queries, such as a KD-Tree [Bentley, 1975], an R-Tree [Guttman, 1984], or any of their variants. The main drawback of using this approach is that high-quality features are slow to calculate. Even though the features of the database light curves are calculated in advance, when a query is performed,

the features of the query light curve have to be extracted, resulting in the main bottleneck of the similarity search process.

We present an experimental analysis with OGLE [Udalski et al., 1996, 2008], MACHO [Alcock et al., 2003, 1997] and KEPLER [Koch et al., 2010] light curve datasets, showing the query times of the Variability Tree performs significantly faster feature extraction approach described in the above paragraph, while retrieving rankings of similar or better quality.

The remainder of this document is organized as follows: Chapter 2 introduces relevant background theory and gives an account of previous work on time series similarity search and light curve analysis. In Chapter 3, the proposed method is described in detail. Chapter 4 describes the datasets that were used on experiments. Chapter 5 gives an account of the software stack we use for implementing the method. Chapter 6 presents experimental results. Finally, in Chapter 7, the main conclusions of the work are presented.

## 2. BACKGROUND THEORY AND RELATED WORK

To understand how the *Variability Tree* works, it is vital to understand the main concepts of Cluster Analysis and to be familiarized with the generic tree data structure. In this section, we briefly explain these topics and how they are important in the context of the *Variability Tree*. Also, we review some essential topics of text retrieval that inspire this work.

### 2.1. Clustering and $k$ -medoids

Clustering is the task of identifying groups of similar objects in a data set, such that elements in the same group are similar to each other, while elements in different groups are dissimilar [Gorunescu, 2011]. To find this groups, called clusters, we need a clustering algorithm and a distance measure. A clustering algorithm is a set of rules to find clusters, while a distance measure is a function that tells us how different are two elements of the data set. For example, if the elements in our data set are vectors in  $\mathbb{R}^n$ , a valid distance measure is the Euclidean distance.

The selection of a specific clustering algorithm affects not only the procedure how the clusters are found but also the final result. Besides that, different clustering algorithms receive different parameters.

For example, the popular  $k$ -means clustering algorithm receives as input the number of clusters  $k$  to be found and retrieves a set of  $k$  centres and the assignment to one cluster for each element of the dataset, such that each of the  $k$  centres correspond to the centroids—or means— of the clusters and each element belongs to the cluster with the closest centre.

In the construction of the *Variability Tree*, we need to apply a clustering algorithm to a set of light curve subsequences. These subsequences are fragments taken from the light curves in the dataset. Subsequences are also light curves themselves, but they have observations in a smaller period of time. For this purpose, we need to choose a similarity a measure that allows us to compare light curves and a clustering algorithm that lets us use such a

similarity measure. The *Vocabulary Tree* in which the *Variability Tree* is inspired uses  $k$ -means for clustering, which works only in Euclidean spaces in its original form. Due to their irregularity, light curves cannot be directly compared using the Euclidean similarity measure, which makes us impossible to use  $k$ -means.

The  $k$ -medoids (or Partition Around Medoids) [Kaufman & Rousseeuw, 1987] clustering algorithm is a modified version of  $k$ -means, which chooses data exemplars as cluster centres, instead of setting the centre on the true mean of the cluster, as  $k$ -means does.

We use  $k$ -medoids because we need to have exemplars in the centres of the clusters, so they can be compared to other exemplar. Also,  $k$ -medoids has the advantage that lets us choose the number of clusters, giving us more control over the structure of the tree, which does not hold for other exemplar-centred clustering algorithms, such as Affinity Propagation [Dueck & Frey, 2007]. The requirement of having exemplar-centred clusters is the same reason we discarded using traditional hierarchical clustering algorithms, such as Hierarchical Agglomerative Clustering.

When using  $k$ -medoids, the pairwise distance between all the exemplars have to be calculated in advance and stored in a distance matrix. The distance matrix and the number of clusters  $k$  are given to the algorithm as parameters.  $k$  medoids identifies  $k$  exemplars as centres and assigns every element to one of  $k$  clusters.

## 2.2. Time Warp Edit Distance

We use the Time Warp Edit Distance (TWED) [Marteau, 2009] as a distance measure for light curve subsequences. The TWED is a measure of distance for time series defined by the number of edit operations needed to transform one time series into another and the cost of these operations. It is designed to compare two time series, regardless of the number of data points each one has, which makes it perfect for irregular time series, such as light curves. It has already been successfully used by Mackenzie et al. [2016] for comparing light curves. A

detailed explanation of TWED calculation can be found in Marteau [2009] and Mackenzie et al. [2016].

### 2.3. Trees

The *Variability Tree* is, as its name indicates, a tree data structure. In this section, we present the main characteristics of trees, as we understand them, together with relevant nomenclature that will be used in further sections.

A tree consists of a series of nodes which are connected by unidirectional links, following certain restrictions. Each node has a unique identifier and a value. The value can be any piece of data, depending on the purposes of the tree.

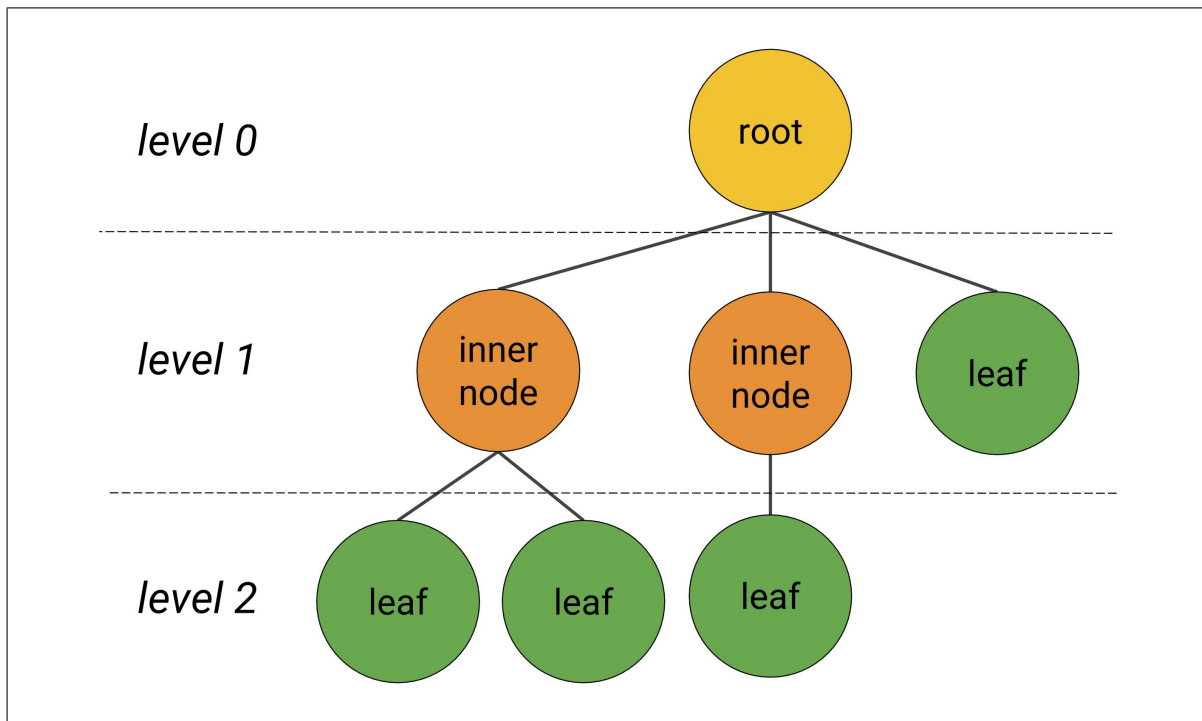


Figure 2.1. The generic tree data structure with its main parts.

If node *A* has a link pointing to node *B*, we say *A* is the *parent* of *B* and *B* is *A*'s *child*. A node in the tree can have any number of children, but only one parent. Every node but one



must have a parent. The node in the tree that doesn't have a parent is called the *root* of the tree and it is an *ancestor* of every other node in the tree. A node in the tree doesn't necessarily have to have children. If a node doesn't have any children, it is called a *leaf* of the tree. On the contrary, if it has children, and it is not the root of the tree, we call it an *inner node*.

We say a node belongs to *level*  $n$ , or has *depth*  $n$ , where  $n$  is the number of links between itself and the root node. In other words, the root node belongs to level 0, its children belong to level 1, and so on. Figure 2.1 shows a visual representation of a tree, displaying its main parts and their relations.

Another concept worth mentioning is the *branching factor* of the node. This is the number of children a node has. Also, if every inner node of the tree has the same branching factor, it is common to say this number is the branching factor of the tree.

## 2.4. Text retrieval review

Although our direct inspiration comes from image retrieval, the main concepts of the Variability Tree have their origins on text retrieval. In this section, we will review key concepts of text mining and how they give origin to the *Variability Tree*.

When building a text retrieval framework, the first thing to be done is to create a numeric representation for the documents in the data set. There are several ways transform documents into numeric vectors, but, arguably, the most canonical model is Bag of Words (BoW). Typically, before building the representation, some preprocessing steps are performed to the text, but these are not relevant for our purposes.

To transform documents into vectors using BoW, first, a dictionary is built by assigning a unique identifier to each of the words that appear in the set of documents. Then, each document in the dataset is represented by a vector, that indicates the *time frequency* of each word—that is, how many times each word of the dictionary was mentioned in it. Usually,

the identifier of the word serves as the index and the number of repetitions is stored in each component—.

More complex representations use weighting schemes, that assign a value of importance to each word of the dictionary. One of the most widely used weighting schemes is *tf-idf* [Salton & McGill, 1983]. This scheme is based on the notion that an infrequent word gives more information than a frequent one. For each word, we calculate a weight called *inverse document frequency* (*idf*). This weight depends on the inverse of the number of documents the word appears. The *idf* of a word multiplied by the time frequency (*tf*) of that word in a document results in the *tf-idf*. To represent a document as a vector using this scheme, we replace the number of occurrences of each word in the document for the *tf-idf* of the word in the document.

$$\text{idf}(\text{word}) = \ln \frac{N}{|\{d \in D : \text{word} \in d\}|} \quad (2.1)$$

$$\text{tf-idf}(\text{word}, \text{doc}) = \text{idf}(\text{word}) \times \text{tf}(\text{word}, \text{doc}) \quad (2.2)$$

A main component of text retrieval engines are inverted indices. Inverted indices—also known as inverted files—work like a glossary of terms, where you can look up a word and be directed to the places where it was mentioned. An inverted index is a data structure that stores for each word in the dictionary, the list of documents where it was mentioned, together with the number of occurrences of the word in the documents. This makes possible to answer the question of which documents mention a certain word  $w$ , without having to search document by document if  $w$  is present.

With documents represented as vectors and having built inverted files, we can perform text relevance queries. A query consists on a phrase and we expect the text retrieval engine to retrieve us a list of documents from the dataset sorted by relevance. To perform a query, first, the query is represented by a factor, using *BoW* or *tf-idf*, the same way as documents in

Table 2.1. Text retrieval concepts exemplified.

(a) A document dataset.

- D1** A red car.  
**D2** A red apple and a red balloon.  
**D3** A yellow banana and a green apple.

(b) A vocabulary extracted from the document set in (a).

ID	word	IDF
1	a	0.0
2	and	1.1
3	apple	0.4
4	balloon	0.5
5	banana	1.1
6	car	1.1
7	green	1.1
8	red	0.4
9	yellow	1

(c) A BoW representation for the document set in (a).

	1	2	3	4	5	6	7	8	9
<b>D1</b>	1	0	0	0	0	1	0	1	0
<b>D2</b>	1	0	1	0	1	0	0	2	0
<b>D3</b>	2	1	1	1	0	0	1	0	1

(d) A TF-IDF representation for the document set in (a).

	1	2	3	4	5	6	7	8	9
<b>D1</b>	0	0	0	0	0	1.1	0	0.4	0
<b>D2</b>	0	0	0.4	0	0	0	0	0.8	0
<b>D3</b>	0	1.1	0.4	0.4	1.1	0	1.1	0	1.1

the dataset. Then, inverted files are used to find which documents contain the words in the query phrase. Finally, documents are sorted according to their similarity to the query vector. Usually vectors are previously normalized, and the angle between them is used as a measure of similarity.

## 2.5. From the Vocabulary Tree to the Variability Tree

The approach above was shown to be useful not only for text retrieval. Sivic et al. [2003] built an image retrieval framework using the same approach. In their work, the authors propose an image representation as a document composed of *visual words*. To do this, they extract features from images and build a dictionary of *visual words* using  $k$ -means to cluster the features. In this framework, each cluster represents a *visual word*. An occurrence of a *visual word*  $w$  in a document  $d$  is when a feature extracted from  $d$  belongs to the cluster  $w$ . Based on this, they extrapolate to images the approach described above, using *tf-idf* for vector representation and inverted files for fast retrieval.

In Nister & Stewenius [2006], the authors present the *Vocabulary Tree*, an image recognition framework that uses the text retrieval approach that was brought to computer vision by Sivic et al. [2003]. They also build a dictionary of visual words, but instead of performing  $k$ -means once, they build a tree structure by dividing clusters into sub-clusters repeatedly, where each sub-cluster of any of the clustering levels represents a visual word. This results in a larger vocabulary and is shown to improve retrieval results.

The *Variabilty Tree* extends the *Vocabulary Tree* approach to light curves. In the following section we explain how the tree is built by clustering subsequences of light curves, allowing fast light curve retrieval.

## 2.6. Related Work

Although similarity search in time series has been widely studied, as mentioned before, to our knowledge, there are no published works that present a method specifically designed for light curves. Agrawal et al. [1993] were the first using the Discrete Fourier Transform (DFT) for time series similarity search. Agrawal et al. [1995] proposed an indexing model for time series based on the Dynamic Time Warping (DTW) distance measure using an R\*-Tree. Chan & Fu [1999] introduced for the first time the use of the Discrete Wavelet Transform (DWT) as a replacement for DFT. Keogh et al. [2001] presented Piecewise Alternative Approximation,

an alternative to DFT and DWT. Popivanov & Miller [2002] introduced the use of smooth wavelets in DWT time series similarity search. Kontaki & Papadopoulos [2004] proposed an index for similarity search in time series with streaming data. Sart et al. [2010] used GPUs and FPGAs to accelerate subsequence search using the Dynamic Time Warping similarity measure. Rakthanmanon et al. [2013] proposed a series of optimizations to existing methods, which led them to perform queries on massive datasets of up to a trillion elements. Baydogan & Runger [2016] introduced a new representation for time series, based on an autoregressive kernel, along with a similarity measure for this representation.

All the methods listed above assume evenly sampled data, either in the transformations or in the similarity measures they use. Therefore, they cannot be directly applied to light curves. To do this, light curves would have to be either interpolated and re-sampled, or binned. This could lead in many cases to bad approximations, since light curves are often poorly sampled and have gaps with no information.

Several works present feature extraction methods on light curves, mostly based on statistics or periodogram analysis [Debosscher et al., 2007, Huijse et al., 2012, Kim et al., 2009, 2014, 2011, Richards et al., 2011], but as mentioned previously, highly descriptive features are often computationally expensive.

Our work is inspired on Nister & Stewenius [2006], where the authors propose using  $k$ -means to build a hierarchy of *visual words*, together with a tree data structure to perform efficient similarity queries for image retrieval. It also relies heavily on the idea of clustering of subsequences proposed by Mackenzie et al. [2016] for light curve feature extraction. Specifically, we borrow from the latter the ideas of subsequence sampling and using the Time Warp Edit Distance (TWED) [Marteau, 2009] as a distance metric for light curves.

### 3. DESCRIPTION OF THE METHOD

The *Variability Tree* is a data structure that allows performing fast similarity search and retrieval from light curve data sets. In this section, we describe in detail how the *Variability Tree* is built from the light curve data set and present an algorithm to perform queries over the tree.

It is important to note that the construction of the data structure is an *off-line* process. This means it is done in advance to the queries and does not have to be repeated each time a query is performed.

The construction of the *Variability Tree* consists of five steps. (i) First, we standardise the amplitude of the light curves in the dataset. (ii) In the second step, we take from the data set a random sample of fragments of time series called subsequences. (iii) In the third step, these subsequences are hierarchically clustered with  $k$ -medoids, creating a tree structure. (iv) The fourth step consists in running a sliding window over every time series in the dataset and adding the resultant subsequences to the tree that has been previously built. (v) Finally, in the fifth step, we build a matrix representation for the dataset, based on the count of their subsequences in the tree nodes. These five steps are described in Sections 3.1.1, 3.1.2, 3.1.3, 3.1.4 and 3.1.5 respectively.

The algorithm to query the *Variability Tree* has five main steps. (i) The first step consists in standardising the query light curve and extracting subsequences from it, as it was done before with the light curve in the dataset. (ii) The second step consists in adding the extracted subsequences to the tree. (iii) In the third step, a vector representation of the query is built, based on the number of subsequences that reached each tree node. (iv) The fourth step consists in creating a reduced version of the query vector and the matrix representation of the dataset. This step is the key point of this method to achieve fast results. (v) Finally, in the fifth step, we calculate the distances between the reduced query vector and the rows of the query matrix—representing the light curve in the data set—, and these distances are sorted to

retrieve the result in the form of a ranking. These five steps are described in Sections 3.2.1, 3.2.2, 3.2.3, 3.2.4, 3.2.5.

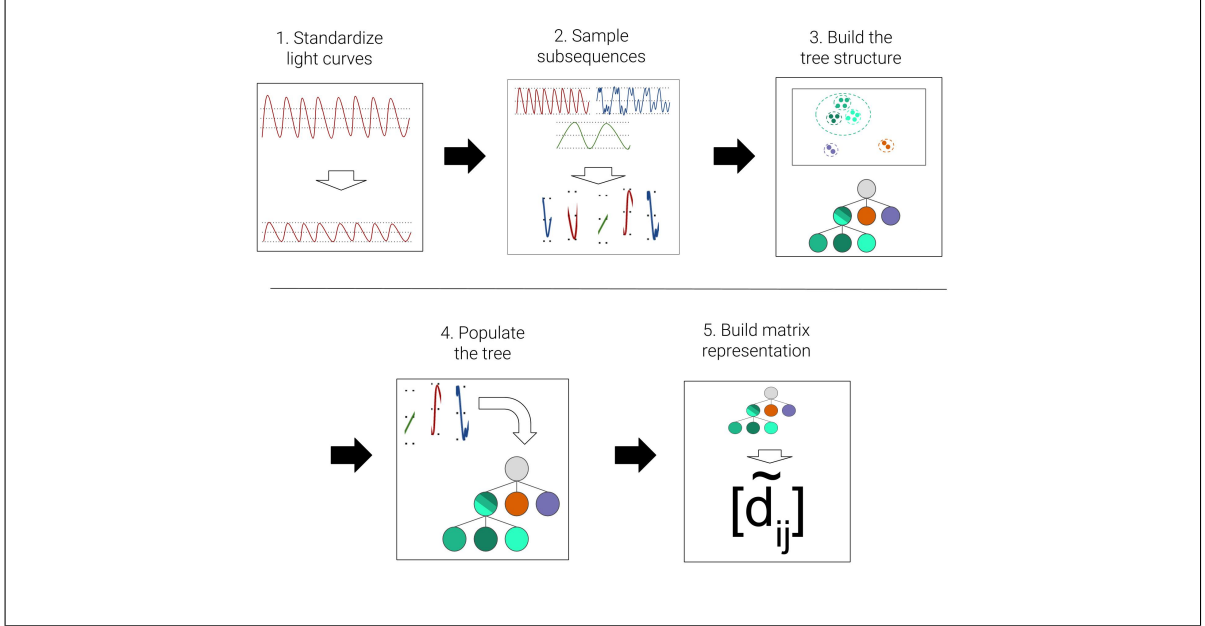


Figure 3.1. The *Variability Tree* building process.

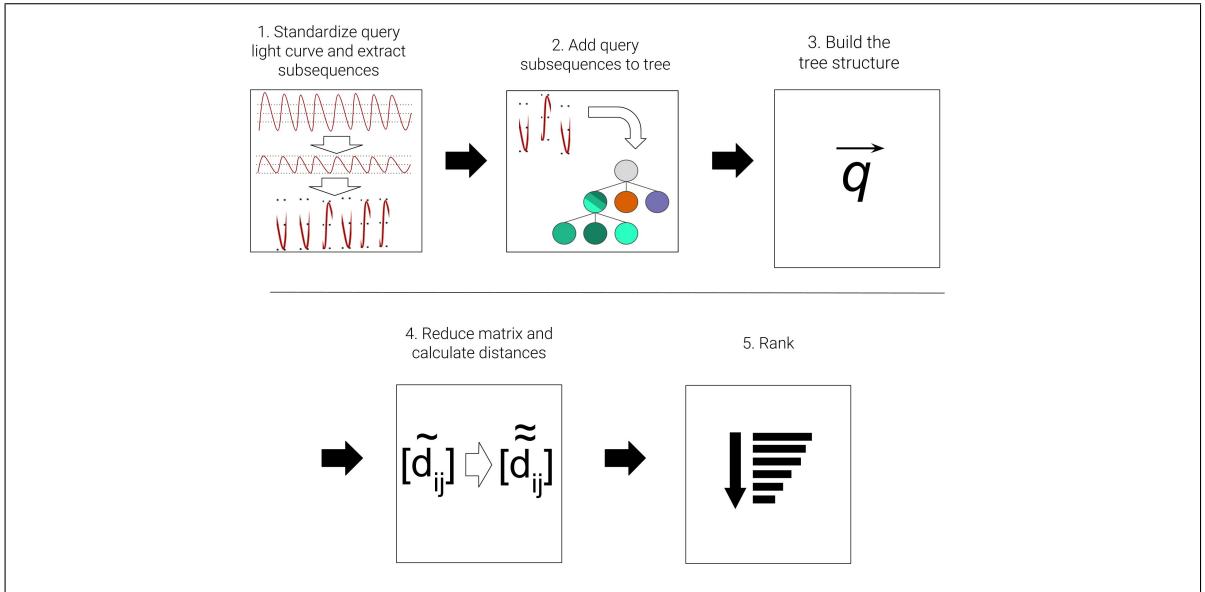


Figure 3.2. The *Variability Tree* querying process.

### 3.1. Building the *Variability Tree*

To build the *Variability tree*, we start with a data set of light curves that are in the same band, but can have different lengths and represent objects of different nature.

#### 3.1.1. Standardise light curves

The first step to build the *Variability Tree* is to standardise the magnitudes of all the light curves in the dataset. We subtract their mean magnitude and divide by their standard deviation, leaving every light curve with a mean magnitude of 0 and standard deviation of 1. Despite losing information about the absolute magnitude of the light curves, we found that better results were achieved with standardised light curves.

#### 3.1.2. Sample subsequences

In this step, we take a random sample of  $N$  light curves from the data set. From each of the sampled light curves, we extract a fragment from a time window of size  $t_w$ . The starting point of this time window is taken randomly. We call these fragments *subsequences*.  $N$  and  $t_w$  are parameters of the model, so their values have to be given by the user.  $N$  must be lower than the total number of light curves in the data set.

Since the light curves are not uniformly sampled, the extracted subsequences do not necessarily have the same number of data points. The parameter  $t_w$  defines the size of the window in terms of time and not in terms of the number of observations.

In subsequent steps, we will need to compare these subsequences using the Time Warp Edit Distance (TWED) measure (explained in Section 2.2). In order to do so, we shift the time values of the extracted subsequences, setting their initial values to zero.

If we do not have labels for the dataset, these  $N$  light curves are uniformly sampled from the data set, leading to a completely unsupervised model. On the other hand, if the data set is labelled, we can optionally perform a stratified sampling. This means that we sample the light



curves in a way that the size of the sample taken from each class is proportional to the total number of light curves in it. The goal here is to build a balanced sample, that is representative of the class distribution in the dataset.

### 3.1.3. Build the tree structure

In this step, we use the  $k$ -medoids clustering algorithm to build a tree data structure from the previously obtained subsequences. To do this, we calculate the pairwise TWED distance between the extracted subsequences and store them in a matrix.

After doing this, we run  $k$ -medoids to cluster the subsequences. The algorithm receives the distance matrix and the number of clusters  $k$  as input.

As a result, this algorithm identifies  $k$  of the  $N$  subsequences as cluster centres and each of the  $N$  subsequences as part of one cluster. The clustering algorithm is run again over each of the obtained clusters, creating subclusters. The process of subdividing clusters into smaller clusters is repeated  $m$  times, dividing each subcluster into  $k$  smaller subclusters. This is applied to every subcluster unless it has less than  $k$  elements and cannot be subdivided.

This resulting hierarchy of  $m$  levels of clusters and subclusters serves us to build the structure of the *Variability Tree*. In this tree, each of the clusters and subclusters is represented by a node. When a cluster is divided into subclusters, these are represented by the children of its respective node.

In each node of the tree, we store the centre of its respective cluster—that is, the subsequence which is at its centre—. Figure 3.3 illustrates how a tree is built from a hierarchy of clusters.

Both  $k$  and  $m$  are parameters of the *Variability Tree* and, therefore, have to be defined by the user. These parameters determine the number of nodes it will have. The parameter  $k$ , which is the number of clusters in  $k$ -medoids, becomes the *branching factor* of the tree;




Figure 3.3. Tree construction through top-down hierarchical  $k$ -medoids. For simplicity, light curves are represented as points in a 2-dimensional space. In this example, branching factor  $k$  is set to 3 and maximum level  $m$  is set to 2, leading to a maximum possible of  $\sum_{i=0}^2 3^i = 13$  nodes. Since clusters at nodes  $n_2$  and  $n_3$  have less than  $k$  elements, they do not split into sub-clusters, resulting a total of 7 nodes. Clusters at nodes  $n_4$ ,  $n_5$  and  $n_6$  do not split because they belong to level 2, which is set as maximum.

while  $m$  defines *maximum depth* of the tree. A tree with branching factor  $k$  and maximum depth  $m$  will have a maximum number of  $\sum_{i=0}^m k^i$  nodes.

The idea behind clustering subsequences is to find subsequences that could represent building blocks for light curves, the same as words are the building blocks of documents (see Section 2.4). This process is analogous to what is done in image retrieval by clustering previously extracted features to find the best *visual words* to describe them.

#### 3.1.4. Populate the tree

Up to this point, we have built the structure of the *Variability Tree*. Now it has to be populated with all the light curves in the data set. This is done by extracting a series of subsequences from each of the light curves in the data set and adding these subsequences to the tree.

To extract subsequences from a light curve, we run a sliding window of time-step  $t_s$  and time-window  $t_w$  across it. This sliding window starts at the beginning of the light curve, taking a subsequence that contains all data points between its first observation, at  $t_0$  and  $t_0 + t_w$ . Then, it slides across the light curve and extracts a subsequence between the time of first observation after  $t_s$ , let's say  $t'_s$ , and  $t'_s + t_w$ . The window keeps sliding by a time-step of size  $t_s$  and extracting subsequences until it reaches the end of the light curve. This means the number of subsequences extracted for each light curve depends on the length of the light curve.

The parameter  $t_w$  is the same that was mentioned in Section 3.1.2. The value of  $t_s$  is also a parameter of the model; and, as well as  $t_w$ , is fixed in time and not in number of observations. Usually,  $t_s$  should be much smaller than  $t_w$  so that the sliding window produces overlapping subsequences. This is done to avoid missing any patterns in the light curve.




Figure 3.4. Subsequences are extracted from every light curve in the data set using a sliding window.

Subsequences are added into the tree through its root and they are propagated down through the node representing their closest cluster centre until they reach a leaf. This is done

in the following manner: first, we calculate the *TWED* distances between the subsequence we want to add and each of the subsequences stored in the first level of the tree—i. e. the centres of the clusters found by  $k$ -medoids on its first execution—. After this, we select the node whose centre has the smallest distance to the subsequence. If this node is not a leaf, we repeat the process and calculate the distances between the subsequence and the subsequences stored in the children of the selected node. We repeat this process consecutively until the added subsequence reaches a leaf.

In each leaf node of the tree, we keep a record of how many subsequences of each light curve have reached that leaf. Initially, all these records are empty. Also, non-leaf nodes have records that show how many subsequences of each light curve have passed through them. Records in non-leaf nodes are virtual and not stored in memory because they can be easily obtained by adding up the records of their children nodes. This makes the *Variability Tree* use much less memory than it would if all registers were physically stored. Figure 3.5 shows how a *Variability Tree* looks after being populated with a dataset of three light curves.

These records work as inverted indices. As explained in Section 2.4, inverted indices allow fast retrieval by storing for each of the words in the dictionary, the list of documents that mention that word and the number of occurrences. What we are doing with this records is exactly the same if we consider that the subsequences stored in the nodes of the tree represent the building blocks of light curves, the same as words are the building blocks of documents.

### 3.1.5. Build matrix representation

Once the *Variability Tree* has been built and populated, we store the information gathered in the data structure in a matrix, which will help us on the querying stage that will be explained further on.

We build a matrix  $D$  of  $l \times n$  dimensions, where  $l$  is the number of light curves in the original database and  $n$  is the number of nodes in the tree. In the position  $D_{ij}$  of the matrix, we store the number of light curve subsequences  $i$  that passed through node  $j$ . In other words,

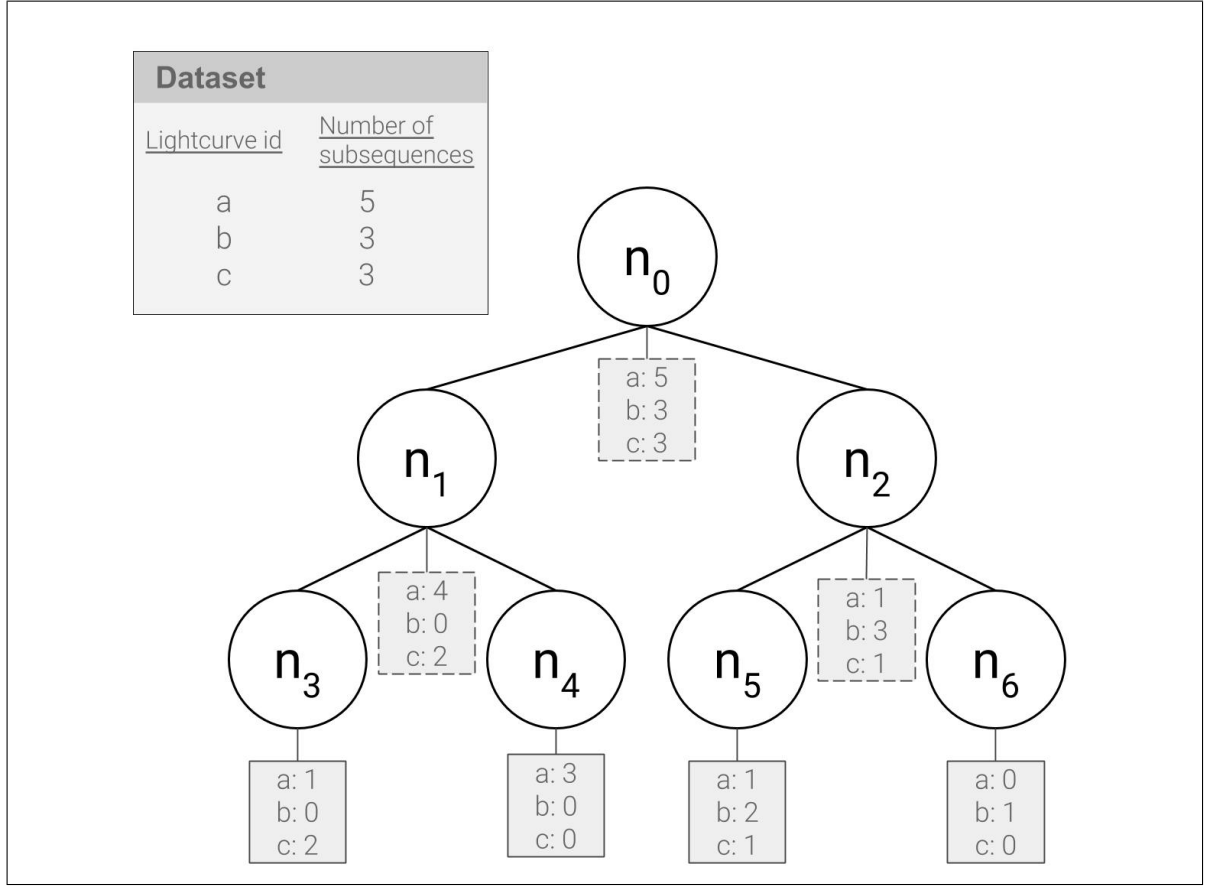


Figure 3.5. Variability Tree after being populated with full dataset. The number of subsequences taken from each light curve is not necessarily the same and depends on the length of the light curve. Each leaf node ( $n_3$ ,  $n_4$ ,  $n_5$  and  $n_6$ ) stores a register of how many subsequences of each lightcurves reached there. For example,  $n_3$  was reached by 1 subsequence of light curve **a**, 2 of light curve **c** and none of light curve **b**. Non-leaf nodes have virtual dictionaries that are calculated by adding up the values of their children’s dictionaries.

each light curve in the original dataset is represented by a row in matrix  $D$ , indicating how many of its subsequences passed through each of the nodes of the *Variability Tree*. This is a sparse matrix, meaning that many of its positions have zeroes. When implementing the method, a sparse matrix representation has to be used, since traditional matrix representations will use unnecessary amounts of space to store null positions.

In this step, we also calculate a weight for each node in the tree. As explained in 2.4, the *idf* weighting scheme assigns more weight to a word that appears in fewer documents.

Applying the same principle, we assign more weight to a node that where subsequences of fewer light curves have passed and assign less weight to crowded nodes.

The weight of node  $j$  is calculated by:

$$w_j = \ln \frac{l}{l_j} \quad (3.1)$$

where  $l$  is the total number of light curves in the data set and  $l_j$  is the number of light curves with at least one subsequence that passed through node  $j$ .

We use these weights to obtain a weighted version of the matrix  $D$ . We denote it as  $\bar{D}$  and its elements are defined as follows:

$$\bar{d}_{ij} = w_j d_{ij} \quad (3.2)$$

Finally, the rows in matrix  $\bar{D}$  are divided by its norm. As will be explained later, having rows as unitary vectors will allow us to perform fast distance computations. Also, we don't lose relevant information by doing this, since we are not interested in the absolute count of subsequences in a node, but the relative count—that is, the number of subsequences that passed through one node compared to the number of subsequences that passed through the others—.

### 3.2. Data retrieval: querying the tree

As stated in Chapter 1, the input for a query can be any light curve—either one that is already contained in the dataset or not—and its output is a list of light curves from the data set ranked by its similarity to the input. Now we explain the five steps that are performed by the model, starting from a single target light curve—or query light curve—and ending in a ranking.

### 3.2.1. Standardise query light curve and extract subsequences

The first thing we do with the target light curve is to standardise its amplitude by dividing it by its mean magnitude and dividing by its standard deviation. After that, we extract subsequences from the target light curve using a sliding window of time-step  $t_s$  and time-window  $t_w$ . These two procedures are identical to what was done with the data set light curves and explained respectively in Sections 3.1.1 and 3.1.4.

### 3.2.2. Add query subsequences to tree

The subsequences that we have just extracted are added to the tree until each reaches a leaf, the same way it was previously done with the light curves in the database, as explained in Section 3.1.4. In each node, we keep a record of how many subsequences of the query light curves passed through it. This record is kept separately from the one that stores the number of subsequences from the data set light curves.

### 3.2.3. Build query vector representation

After adding to the tree the subsequences extracted from the query light curve, we copy to a vector  $\mathbf{q}$  the query data that we recorded in the tree.

This vector stores how many subsequences of the query light curve passed through each node of the tree, and it is weighted in the same manner as rows in matrix  $\bar{D}$ . Its components are defined by the following relation:

$$q_j = w_j p_j \tag{3.3}$$

where  $w_j$  is the weight of node  $j$  and  $p_j$  is the number of subsequences of the query light curve that passed through node  $j$ . After this, the query vector  $\mathbf{q}$  is normalized, as it was done with the rows of the matrix  $\bar{D}$ .

### 3.2.4. Reduce matrix and calculate similarity

Now that we have the data set light curves represented as the rows of matrix  $\bar{D}$  and the query light curve represented as vector  $\mathbf{q}$ , we can calculate a similarity measure between them. To do so, we use the cosine similarity, which is used in the text retrieval approaches we are being inspired by [Nister & Stewenius, 2006, Sivic et al., 2003]. The cosine similarity between two vectors is defined as the cosine of the angle between them. That is,

$$\text{similaridad}(\mathbf{a}, \mathbf{b}) = \frac{\mathbf{a} \cdot \mathbf{b}}{\|\mathbf{a}\| \|\mathbf{b}\|} = \cos(\theta) \quad (3.4)$$

where  $\theta$  is the angle between  $\mathbf{b}$  and  $\mathbf{q}$ .

Since the rows of matrix  $\bar{D}$  and the query vector  $\mathbf{q}$  are already normalized, the calculation of the cosine similarity is reduced to the calculation of the dot product between them. In other words,  $\bar{D}\mathbf{q}$  will result in a vector of length  $l$ , with the similarity between the query vector and each of the light curves in the data set.

Although this operation gives us effectively a similarity measure between the query and every light curve in the data set, it can be improved by avoiding to do some unnecessary calculations. Usually, there will be several light curves in the data set that do not share any node in the tree with the query. In these cases, the dot product between  $\mathbf{q}$  and its respective row will be zero. Since we only want to get similar light curves, these light curves are not relevant in our resulting ranking. Using the inverted indices we kept in the leaves of the trees, we can select from matrix  $\bar{D}$  only the rows that correspond to light curves that share at least one node with the query. We can build a new matrix selecting only the relevant rows.

Our  $\bar{D}$  matrix has also columns that can be ignored to calculate the dot product. Usually, there will be nodes where none of the subsequences from the query light curve passed. The columns in  $\bar{D}$  representing these nodes can be ignored since their values will be multiplied by zeroes when calculating their cosine similarity to  $\mathbf{q}$ .



Taking this into consideration, we build a reduced matrix  $\bar{\bar{D}}$  by selecting from  $\bar{D}$  only its relevant rows and columns. This can be done by a fast matrix slicing operation. We also build a reduced query vector  $\bar{q}$ , by eliminating its null positions—that is, the positions representing nodes where no query subsequences from the passed—.

Now that we have removed unnecessary data, we calculate  $\bar{\bar{D}}\mathbf{q}$  to obtain the similarity between the query and every light curve in the dataset that has a cosine similarity larger than zero.

### **3.2.5. Calculate distances and sort ranking**

Finally, the ids of the light curves in the data set are sorted according to their similarity to the query. These sorted ids, together with their similarity scores constitute the output of the querying algorithm.

## 4. DATA

For experimental purposes, we used photometric data from OGLE, MACHO and KEPLER surveys. In Chapter 4.1 we present the nature and composition of these datasets, while in Chapter 4.2 we present the composition of query-sets. These query-sets are subsets of light curves we selected to use as targets for performing queries.

### 4.1. Datasets

#### 4.1.1. Ogle III Catalog of Variable Starts (OGLE)

The OGLE Experiment [Udalski et al., 1996] is an ongoing project that started on 1992 aiming to find microlensing events on the Galactic Bulge and the Magellanic Clouds. Its third phase, OGLE-III, surveyed the sky between 2001 and 2009, producing a large amount of labelled photometric data. These data were published as the Ogle III Catalog of Variable Starts (OIII-CVS) [Udalski et al., 2008].

The published light curves are in two bands: I-band and V-band. However, we only used the I-band for our experiments, since light curves in the V-band did not have enough observations for our purposes.

For our experiments, we used 444,503 light curves from OIII-CSV, from which 440,850 were labelled. The class distribution of the labelled subset is shown in Table 4.1.

Table 4.1. Class distribution of OGLE labelled set.

Class Name	Abbreviation	No. of elements
Classical Cepheids	CEP	8026
Type II Cepheids	T2CEP	609
Anomalous Cepheids	ACEP	83
RRLyrae	RRLYR	44262
Long Period Variables	LPV	343816
Double Period Variables	DPV	137
Delta Scuti	DSCT	58
Eclipsing binaries	ECL	43859

Table 4.2. Class distribution of MACHO labelled set.

Class Name	Abbreviation	No. of elements
RRLyrae ab	RRL AB	7405
RRLyrae c	RRL C	1765
RRLyrae e	RRL E	315
Cepheid Fundamental	Ceph Fund.	1185
Cepheid 1st Overtone	Ceph 1st.	683
Long Period Variable WoodA	LPV A	315
Long Period Variable WoodB	LPV B	822
Long Period Variable WoodC	LPV C	1134
Long Period Variable WoodD	LPV D	778
Eclipsing Binary	EB	6835
RR Lyrae and GB blends	RRL + GB	237

#### 4.1.2. The MACHO project

The MACHO project [Alcock et al., 1997] started on 1992 with the objective of finding massive compact halo objects that would explain the presence of dark matter in the Milky Way. As a by-product, it produced more than 80 billion photometric measurements from the Galactic Bulge, the Large Magellanic Cloud and the Small Magellanic Cloud [Alcock et al., 1999]. Light curves were obtained on two bands: blue and red. On 2001, the MACHO Collaboration Team published a catalogue that included labels for 21,474 variable stars in the Magellanic Clouds [Alcock et al., 2003].

For our experiments, we used a set of 19,927,044 light curves from MACHO survey, corresponding to objects on the Large Magellanic Cloud, from which 20,391 were labelled.

The class distribution of the labelled subset is shown in Table 4.2. It is worth noting that in this dataset RRLyrae stars are subdivided into three subcategories and Long Period Variable are subdivided into four subcategories, while these are considered two big categories in our OGLE dataset.

### 4.1.3. The Kepler Mission (KEPLER)

The Kepler mission [Koch et al., 2010] was launched on 2009 by NASA with the main purpose of finding Earth-size exoplanets. As a result, it produced a huge amount of high-resolution photometric measurements.

Kepler light curves were published in two cadences: Short Cadence (SC), which has observations every 58.89 seconds, and, and Long Cadence (LC), which has observations every 29.4 minutes. It is worth noting that both LC and SC data have a much shorter cadence than OGLE and MACHO data. While MACHO has approximately one observation per 2 days and OGLE has around one observation per 6 days, Kepler long cadence data has around 50 observations per day. To make Kepler data more manageable, we downsampled the light curves to get one observation per day. We performed our experiments with these downsampled light curves. We used 207,617 light curves from this dataset. We didn't use any labelled data from this survey.

## 4.2. Query-sets

For each of the datasets mentioned in Chapter 4.1, we built a query-set. A query-set consisted of a set of light curves for which we would perform similarity queries to obtain their most similar light curves on a dataset.

We built the query-sets by taking a random sample of light curves for each of the classes in each dataset. Table 4.3 shows the composition of query-sets for each of the three dataset we used.

Table 4.3. Class distribution of query-sets built for OGLE, MACHO and KEPLER datasets.

(a) OGLE		(b) MACHO		(b) KEPLER	
Class	Count	Class	Count	Class	Count
CEP	100	RRL AB	100	<i>Unlabelled</i>	448
T2CEP	98	RRL C	100		
ACEP	83	RRL E	100		
RRLYR	100	Ceph Fund.	100		
LPV	100	Ceph 1st.	100		
DPV	100	LPV A	100		
DSCT	51	LPV B	100		
ECL	86	LPV C	100		
		LPV D	100		
		EB	100		
		RRL + GB	100		

## 5. IMPLEMENTATION

### 5.1. Software stack

Our implementation of the method was written in Python 3.5 using numpy [Van Der Walt et al., 2011], and numba [Lam et al., 2015] for efficient numerical computation, pandas [McKinney, 2010] for data manipulation and scipy [Oliphant, 2007] for sparse matrix representation. For  $k$ -Medoids, we used Szalkai [2013] implementation and for TWED we used Mackenzie et al. [2016]. Our code is open-source under the MIT License and available at <https://github.com/luvalenz/time-series-vocabulary-tree>.

### 5.2. Parameters

As explained previously in Chapter 3, the *Variability Tree* model has five parameters: the number of subsequences at sampling stage  $N$ , the size of the subsequences to be extracted  $t_w$ , the step of the sliding window  $t_s$ , the branching factor of the tree  $k$  and maximum level of the tree  $m$ . To execute our python implementation, the user has to pass the values of this parameters as arguments to a python script. The values we used for each of the parameters in our experiments are shown in Table 5.1.

For the parameters that define how subsequences are extracted—that is,  $N$ ,  $t_w$  and  $t_s$ —, we chose to use the same values as Mackenzie et al. [2016]. Although the method presented on their paper serves a different purpose than ours, both rely on clustering subsequences of light curves to build a new representation of the data. This led us to assume that the values

Table 5.1. Values of parameters used at experiments.

Parameter	Meaning	Value
$N$	Number of subsequences at sampling stage	20000
$t_w$	Size of subsequences in days	250
$t_s$	Step of sliding window in days	10
$k$	Branching factor of the tree	3
$m$	Maximum level of Variability Tree	10

that had proved successful in the subsequence extraction stage, would also be suitable for our work.

For the branching factor  $k$ , a larger value leads to a tree with more nodes per level. Such a tree preserves more information, but has a longer building time and longer query times, because more distance calculations have to be done each time a subsequence is added. On our experiments, we set  $k$  to 3, allowing us to preserve the maximum amount of information while maintaining reasonable query times. With the same purpose of maintaining reasonable query times, we set the maximum level of the tree  $m$  to 10, setting a hard limit to the number of nodes in the tree.

## 6. EXPERIMENTAL RESULTS

In this section, we present the experiments and results that we perform to test our method. We carry out experiments in two different set-ups.

The first set-up (Section 6.1) aims at measuring the quality of the rankings obtained by the *Variability Tree*. To do this, we work under the assumption that the most similar light curves to a given target should belong to the same class. Therefore, a good ranking should be the one that retrieves as many exemplars as possible of the same class in the upper positions. To be able to measure the quality of the rankings, we only use labelled datasets.

The second set-up (Section 6.2) aims at testing how our method performs and scales with larger datasets. Here, we test our method with the full datasets presented in Chapter 4 (containing no labels). We report obtained rankings, together with the query times.

### 6.1. Experimental set-up No. 1: measuring ranking quality

In order to test the quality of the rankings retrieved by our model, we build a *Variability Tree* for each of the labelled dataset presented in Section 4.1. The first tree is built from 20,391 labelled light curves from MACHO dataset, while the second is built from 440,850 labelled light curves from OGLE dataset.

After this process, we use the previously built trees to perform a series of similarity searches (or queries). We perform similarity queries for light curves in the query-sets described in Section 4.2. For each light curve in the OGLE query-set, we obtain a ranked list of their most similar exemplars in the OGLE data set using the OGLE *Variability Tree*. The same process is repeated for MACHO.

The quality of these rankings is measured using **Precision at  $k$  ( $P@k$ )**, which calculates the ratio of relevant results in the top  $k$  positions of a ranking, that is,

$$P@k = \frac{|\{\text{relevant documents in first } k \text{ positions}\}|}{k} \quad (6.1)$$



In this set-up, we compare our results to the rankings obtained when performing nearest neighbour searches over previously transformed data sets using feature extraction methods. Here, we perform feature extraction on OGLE and MACHO labelled data, transforming the light curves into feature vectors. After that, we generate similarity rankings for the light curves in the query-sets by obtaining the closest vectors using the euclidean distance. We performed these nearest neighbour searches using a KD-Tree [Bentley, 1975], which is a data structure specially designed to do this queries in a fast manner.

We use two different feature-sets to perform nearest neighbour searches and compare them to our proposed method. The first one corresponds to the feature set offered by the FATS Python library [Nun et al., 2015]. This is a collection of 64 features used throughout literature. The second corresponds to a set of automatically learned features presented in Mackenzie et al. [2016] (Clustering Based Feature Learning (CBFL)), which is based on clustering of subsequences.

In Table 6.1, we show the P@k score with  $k = 1, 5, 10, 20$  for rankings obtained using the *Variability Tree* (VT), as well as using nearest neighbour search over the FATS and CBFL feature sets. Results show that our method retrieves better quality rankings than nearest neighbour search in the MACHO datasets while obtaining comparable results to CBFL nearest neighbour search in OGLE.

It is worth noting that we intentionally choose subsets of the indexed data as query-sets to test the model. The reason for making these decision is that we know that if the model is correctly implemented, the first element of a ranking must be the query light curve itself. This serves us as a proof of concept of the model implementation. Then, if the model retrieved a ranking where the first element were different from the query, we would immediately detect an error in the model implementation. This is also the reason why P@1 is equal to 1.0 in all the experiments of this set-up. It is important to mention that this decision does not affect in any way the validity of the experiments, since the rankings from the position 2 and on are exactly the same as they would be if the testing data were not included in the indexed data.

Table 6.1. Precisions at k for rankings obtained in experimental set-up No. 1. Values show mean and standard deviation.

		MACHO	OGLE
P@1	VT	$1.00 \pm 0.00$	$1.00 \pm 0.00$
	FATS	$1.00 \pm 0.00$	$1.00 \pm 0.00$
	CBFL	$1.00 \pm 0.00$	$1.00 \pm 0.00$
P@5	VT	$0.57 \pm 0.30$	$0.59 \pm 0.36$
	FATS	$0.46 \pm 0.27$	$0.53 \pm 0.35$
	CBFL	$0.50 \pm 0.28$	$0.61 \pm 0.37$
P@10	VT	$0.50 \pm 0.31$	$0.52 \pm 0.40$
	FATS	$0.39 \pm 0.28$	$0.46 \pm 0.38$
	CBFL	$0.42 \pm 0.28$	$0.54 \pm 0.41$
P@20	VT	$0.46 \pm 0.32$	$0.49 \pm 0.41$
	FATS	$0.34 \pm 0.27$	$0.42 \pm 0.39$
	CBFL	$0.37 \pm 0.27$	$0.50 \pm 0.43$

Table 6.2. Query time (in seconds) for experimental set-up No. 1. Values show mean and standard deviation.

	MACHO	OGLE
VT	$2.37 \pm 1.13$	$2.53 \pm 1.02$
FATS	$80.78 \pm 30.46$	$74.72 \pm 56.07$
CBFL	$9.47 \pm 3.50$	$8.07 \pm 7.79$

While the quality of the obtained rankings can be comparable, the *Variability Tree* clearly outperforms nearest neighbour search in querying time for both datasets (OGLE and MACHO) and both feature sets (FATS and CBFL). Query time for these experiments are shown in Table 6.2.

Another result worthy of mention is that differences in precision between methods are not even amongst classes. In Figures 6.1 and 6.2, the distributions by class of P@10 are shown for OGLE and MACHO, respectively. For example, we can note that in MACHO the *Variability Tree* notoriously outperforms the other methods for the class *Ceph 1st*, while performing similarly to the others for the classes *LPV B* and *LPV C*.

We also find that light curves of some specific classes are likely to be found similar to light curves of other specific classes, while there are pairs of classes whose light curves are




Figure 6.1. P@10 distribution by class for OGLE dataset.




Figure 6.2. P@10 distribution by class for MACHO dataset.

never found similar. To show this, we classify the first five light curves of each ranking with the class of the target light curve. For this purpose, the rest of the ranking positions higher than five are ignored. Figure 6.3 shows confusion matrices obtained from this process. The rows of the matrices represent the actual class of the light curves retrieved in the ranking and the columns represent the class of the class of the target light curves that were used to classify. We can see that in OGLE, light curves classified as classical Cepheids (CEP) are like to be confused with Anomalous Cepheids (ACEP) and, in MACHO, the three different categories of RR-Lyrae— AB, C and E— are likely to be confused with each other.

We also find that light curves of some specific classes are likely to be found similar to light curves of other specific classes, while there are pairs of classes whose light curves are never found similar. To show this, we classify the first five light curves of each ranking with the class of the target light curve. For this purpose, the rest of the ranking positions higher




Figure 6.3. Confusion matrices for OGLE and MACHO datasets.

than five are ignored. Figure 6.3 shows confusion matrices obtained from this process. The rows of the matrices represent the actual class of the light curves retrieved in the ranking and the columns represent the class of the class of the target light curves that were used to classify. We can see that in OGLE, light curves classified as classical Cepheids (CEP) are like to be confused with Anomalous Cepheids (ACEP) and, in MACHO, the three different categories of RR-Lyrae— AB, C and E— are likely to be confused with each other.

## 6.2. Experimental set-up No. 2: Large datasets

In these experiments, we use MACHO, OGLE and KEPLER datasets. For MACHO, we use 19,927,044 light curves, most of which were unlabelled. The few labelled light curves in the dataset correspond to the ones used in set-up No. 1. For OGLE, we use 444,503 light curves, for most of which we have labels. This dataset also correspond to a superset of the dataset used in 6.1. For KEPLER, we use an unlabelled dataset of 207,617 light curves.

Table 6.3. Average time in seconds for queries in experiments for unlabelled datasets.

Dataset	Time [seconds]
MACHO	$147.42 \pm 48.42$
OGLE	$3.62 \pm 1.33$
KEPLER	$2.95 \pm 1.59$

In this set-up, we build a *Variability Tree* for each of these three datasets, using all the light curves mentioned in the above paragraph. After that, we query each tree with its corresponding query-set. For MACHO and OGLE, we use the same query-sets as in set-up No 1. For KEPLER, we use the query-set of 448 unlabelled light curves.

For these experiments, we calculate the query times, which are shown in Table 6.3, and store the obtained rankings. Figure 6.4 shows some examples of the rankings obtained in this set-up. Also, in <https://github.com/luvalenz/time-series-vocabulary-tree>, we make available for download the rankings for all of the queries performed in this experiment set-up, hoping they can be useful for the scientific community. Due to the lack of labels in this case, we cannot give a quality measure of the retrieved rankings, as we did in Section 6.1.

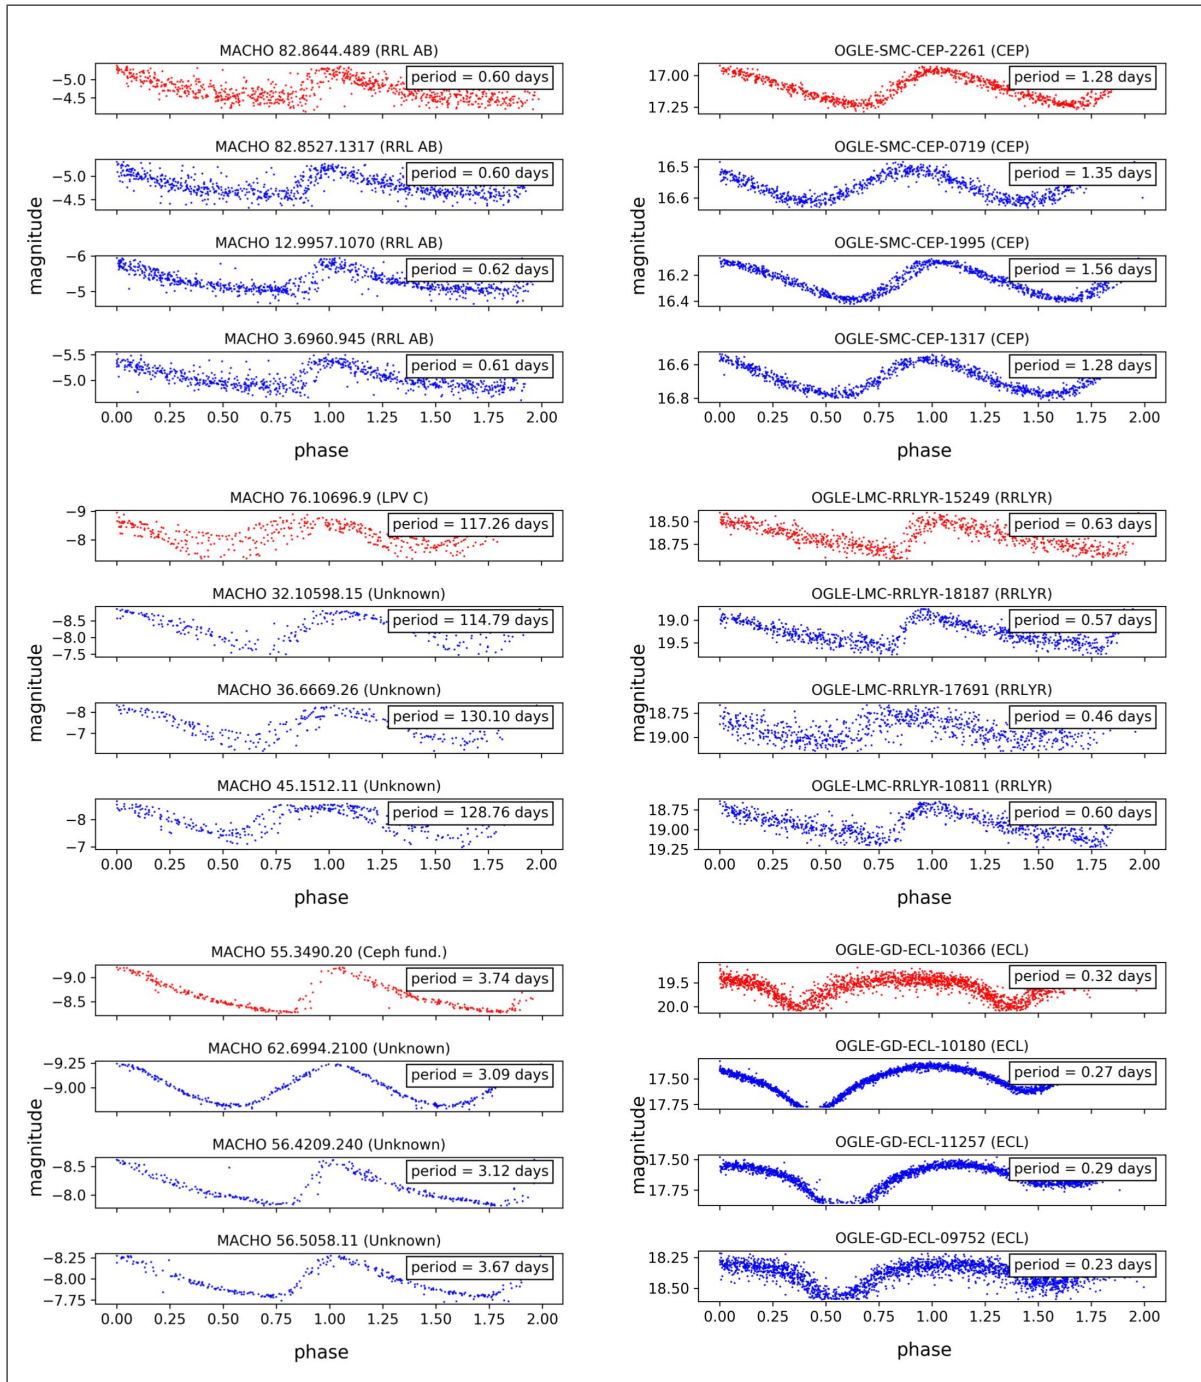


Figure 6.4. Sample queries on MACHO and OGLE data sets. The light curves are shown folded for visualization purposes, although we worked with unfolded raw data.

## 7. CONCLUSIONS

In this work, we have introduced a new method to find similar light curves in large data sets. We presented a novel data structure—the *Variability Tree*—that allows us to create a new compressed representation of light curves that is able to capture the underlying structure of the data. This method integrates feature-extraction and fast search, without having to tackle both problems separately. We have also shown how similarity search can be used to classify light curves without exhaustively training a specialized classifier.

Our results show that our method performs as well or better than nearest neighbour search over feature-vector data, while having a significantly lower computational cost.

We hope this work can aid astronomers to find objects of interest in existing datasets, that due to their size cannot be explored manually. We strongly believe that unsupervised methods, such as this, can help the task of classification, by separating interesting from uninteresting data.

Our Python implementation is open-source and available for download. The values of the parameters shown in Table 5.1 should be adjusted depending on the data.

## REFERENCES

- Agrawal, R., Faloutsos, C., & Swami, A. (1993). Efficient similarity search in sequence databases. *Proceedings of the 4th International Conference on Foundations of Data Organization and Algorithms (FODO '93)*, 8958546, 69–84. Retrieved from [http://link.springer.com/chapter/10.1007/3-540-57301-1\\_{\\_}5](http://link.springer.com/chapter/10.1007/3-540-57301-1_{_}5) doi: 10.1007/3-540-57301-1\_5
- Agrawal, R., Lin, K., Sawhney, H. S., & Shim, K. (1995). Fast similarity search in the presence of noise, scaling, and translation in time-series databases. *Proceedings of the 21st International Conference on Very Large Databases*, 490–501. Retrieved from <http://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.40.4034{&}rep=rep1{&}type=pdf>
- Alcock, C., Allsman, R., Alves, D., Axelrod, T., Bennett, D., Cook, K., ... Welch, D. (2003). *Variable Stars in the Large Magellanic Clouds (MACHO, 2001)* (Vol. 2247).
- Alcock, C., Allsman, R. A., Alves, D., Axelrod, T. S., Becker, A. C., Bennett, D. P., ... Stubbs, C. W. (1997, sep). The MACHO Project Large Magellanic Cloud Microlensing Results from the First Two Years and the Nature of the Galactic Dark Halo. *The Astrophysical Journal*, 486(1981), 697–726. Retrieved from <http://stacks.iop.org/0004-637X/486/i=2/a=697> doi: 10.1086/304535
- Alcock, C., Allsman, R. a., Alves, D. R., Axelrod, T. S., Becker, a. C., Bennett, D. P., ... Welch, D. L. (1999, dec). Calibration of the MACHO Photometry Database. *Publications of the Astronomical Society of the Pacific*, 111(766), 47. Retrieved from <http://arxiv.org/abs/astro-ph/9909227> doi: 10.1086/316469
- Angeloni, R., Contreras Ramos, R., Catelan, M., Dkny, I., Gran, F., Alonso-Garca, J., ... Whitelock, P. (2014). The VVV Templates Project Towards an automated classification of VVV light-curves. I. Building a database of stellar variability in the near-infrared. *Astronomy and Astrophysics*, 567(34), A100. Retrieved from <http://adsabs.harvard.edu/abs/2014A{ }26A...567A.100A> doi: 10.1051/0004-6361/201423904



- Baydogan, M. G., & Runger, G. (2016, mar). Time series representation and similarity based on local autopatterns. *Data Mining and Knowledge Discovery*, 30(2), 476–509. Retrieved from <http://link.springer.com/10.1007/s10618-015-0425-y> doi: 10.1007/s10618-015-0425-y
- Bentley, J. L. (1975, sep). Multidimensional binary search trees used for associative searching. *Communications of the ACM*, 18(9), 509–517. Retrieved from <http://portal.acm.org/citation.cfm?doid=361002.361007> doi: 10.1145/361002.361007
- Bhamidipaty, A., Janakiraman, A., Sarawagi, S., & Haritsa, J. (2002). Building Classifiers With Unrepresentative Training Instances : Experiences From The KDD Cup 2001 Competition. *Workshop on Data Mining Lessons Learnt held in conjunction with the International Conference on Machine Learning*. Retrieved from <http://dsl.cds.iisc.ac.in/pub/KDDCup.pdf>
- Bloom, J. S., Richards, J. W., Nugent, P. E., Quimby, R. M., Kasliwal, M. M., Starr, D. L., ... Law, N. (2012, nov). Automating Discovery and Classification of Transients and Variable Stars in the Synoptic Survey Era. *Publications of the Astronomical Society of the Pacific*, 124(921), 1175–1196. Retrieved from <http://www.jstor.org/stable/10.1086/668468> doi: 10.1086/668468
- Cabrera-Vives, G., Reyes, I., Förster, F., Estévez, P. A., & Maureira, J.-C. (2017, feb). Deep-HiTS: Rotation Invariant Convolutional Neural Network for Transient Detection. *The Astrophysical Journal*, 836(1), 97. Retrieved from <http://stacks.iop.org/0004-637X/836/i=1/a=97?key=crossref.954f0c5097bbe657933c12495ccc349b><http://arxiv.org/abs/1701.00458><http://dx.doi.org/10.3847/1538-4357/836/1/97> doi: 10.3847/1538-4357/836/1/97
- Chan, K.-P. C. K.-P., & Fu, A. W.-C. F. A. W.-C. (1999). Efficient time series matching by wavelets. *Proceedings 15th International Conference on Data Engineering (Cat. No.99CB36337)*, 0–7. Retrieved from <http://ieeexplore.ieee.org/document/754915/> doi: 10.1109/ICDE.1999.754915
- Debosscher, J., Sarro, L. M., Aerts, C., Cuypers, J., Vandenbussche, B., Garrido, R., &

- Solano, E. (2007, dec). Automated supervised classification of variable stars. *Astronomy and Astrophysics*, 475(3), 1159–1183. Retrieved from <http://www.aanda.org/10.1051/0004-6361:20077638> doi: 10.1051/0004-6361:20077638
- Dueck, D., & Frey, B. J. (2007). Clustering by Passing Messages Between Data Points. *Science*, 315(5814), 972–976. doi: 10.1126/science.1136800
- Elorrieta, F., Eyheramendy, S., Jordán, A., Dékány, I., Catelan, M., Angeloni, R., ... Minniti, D. (2016). A machine learned classifier for RR Lyrae in the VVV survey. *Astronomy & Astrophysics*, 595, A82. Retrieved from <https://www.aanda.org/articles/aa/pdf/2016/11/aa28700-16.pdf><http://www.aanda.org/10.1051/0004-6361/201628700> doi: 10.1051/0004-6361/201628700
- Förster, F., Maureira, J. C., San Martín, J., Hamuy, M., Martínez, J., Huijse, P., ... Vera, E. (2016, nov). the High Cadence Transient Survey (Hits) I. Survey Design and Supernova Shock Breakout Constraints. *Not to appear in Nonlearned J*, 45(2), 1–22. Retrieved from <http://stacks.iop.org/0004-637X/832/i=2/a=155?key=crossref.299d3cb682af2367b2752c1623d36cfa><http://dx.doi.org/10.3847/0004-637X/832/2/155> doi: 10.3847/0004-637X/832/2/155
- Gorunescu, F. (2011). *Data Mining: Concepts and Techniques* (Vol. 12). Retrieved from <http://link.springer.com/10.1007/978-3-642-19721-5> doi: 10.1007/978-3-642-19721-5
- Guttman, A. (1984). R-trees. *Proceedings of the 1984 ACM SIGMOD international conference on Management of data - SIGMOD '84*, 47. Retrieved from <http://portal.acm.org/citation.cfm?doid=602259.602266> doi: 10.1145/602259.602266
- Huijse, P., Estevez, P. A., Protopapas, P., Principe, J. C., & Zegers, P. (2014, aug). Computational intelligence challenges and applications on large-scale astronomical time series databases. *IEEE Computational Intelligence Magazine*, 9(3), 27–39. Retrieved from <http://ieeexplore.ieee.org/document/6853469/> doi: 10.1109/MCI.2014.2326100

- Huijse, P., Estevez, P. A., Protopapas, P., Zegers, P., & Pražnička, J. C. (2012, oct). An information theoretic algorithm for finding periodicities in stellar light curves. *IEEE Transactions on Signal Processing*, 60(10), 5135–5145. Retrieved from <http://ieeexplore.ieee.org/document/6215064/> doi: 10.1109/TSP.2012.2204260
- Ivezic, Z., Tyson, J. A., Abel, B., Acosta, E., Allsman, R., AlSayyad, Y., ... Collaboration, f. t. L. (2008, may). LSST: from Science Drivers to Reference Design and Anticipated Data Products. Retrieved from <http://arxiv.org/abs/0805.2366>
- Kaufman, L., & Rousseeuw, P. J. (1987). *Clustering by means of medoids*. North-Holland. Retrieved from <http://ci.nii.ac.jp/naid/10018044334/>
- Keogh, E., Chakrabarti, K., Pazzani, M., & Mehrotra, S. (2001, aug). Dimensionality Reduction for Fast Similarity Search in Large Time Series Databases. *Knowl. Inf. Syst.*, 3(3), 263–286. Retrieved from <http://link.springer.com/10.1007/PL00011669> doi: 10.1007/PL00011669
- Kim, D. W., Protopapas, P., Alcock, C., Byun, Y. I., & Bianco, F. B. (2009, jul). Detrending time series for astronomical variability surveys. *Monthly Notices of the Royal Astronomical Society*, 397(1), 558–568. Retrieved from <http://mnras.oxfordjournals.org/cgi/doi/10.1111/j.1365-2966.2009.14967.x> doi: 10.1111/j.1365-2966.2009.14967.x
- Kim, D.-W., Protopapas, P., Bailer-Jones, C. A. L., Byun, Y.-I., Chang, S.-W., Marquette, J.-B., & Shin, M.-S. (2014, jun). The EPOCH Project. *Astronomy & Astrophysics*, 566, A43. Retrieved from <http://arxiv.org/abs/1403.6131> doi: 10.1051/0004-6361/201323252
- Kim, D.-W., Protopapas, P., Byun, Y.-I., Alcock, C., Khardon, R., Trichas, M., ... Zhang, Y., Y. & Zhao (2011, jul). Quasi-Stellar Object Selection Algorithm Using Time Variability and Machine Learning: Selection of 1620 Quasi-Stellar Object Candidates From Macho Large Magellanic Cloud Database. *The Astrophysical Journal*, 735(2), 68. Retrieved from <http://stacks.iop.org/0004-637X/735/i=2/a=68?key=crossref.175979811d853df1d77fc13a280e3787> doi: 10.1088/0004-637X/735/2/68

- Koch, D. G., Borucki, W. J., Basri, G., Batalha, N. M., Brown, T. M., Caldwell, D., ... Wu, H. (2010). Kepler Mission Design, Realized Photometric Performance, and Early Science. *The Astrophysical Journal Letters*, 713(2), 79–86. Retrieved from <http://iopscience.iop.org/article/10.1088/2041-8205/713/2/L79> doi: 10.1088/2041-8205/713/2/L79
- Kontaki, M., & Papadopoulos, A. (2004). Efficient similarity search in streaming time sequences. *Proceedings. 16th International Conference on Scientific and Statistical Database Management, 2004.*, 63–72. Retrieved from <http://ieeexplore.ieee.org/lpdocs/epic03/wrapper.htm?arnumber=1311194> doi: 10.1109/SSDM.2004.1311194
- Lam, S. K., Pitrou, A., & Seibert, S. (2015). Numba. *Proceedings of the Second Workshop on the LLVM Compiler Infrastructure in HPC - LLVM '15*, 1–6. Retrieved from <http://dl.acm.org/citation.cfm?doid=2833157.2833162> doi: 10.1145/2833157.2833162
- Mackenzie, C., Pichara, K., & Protopapas, P. (2016, mar). Clustering-Based Feature Learning on Variable Stars. *The Astrophysical Journal*, 820(2), 138. Retrieved from <http://stacks.iop.org/0004-637X/820/i=2/a=138?key=crossref.ec69177bb4fc133df178cbd6a800986d> doi: 10.3847/0004-637X/820/2/138
- Marteau, P. F. (2009, feb). Time warp edit distance with stiffness adjustment for time series matching. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 31(2), 306–318. Retrieved from <http://ieeexplore.ieee.org/document/4479483/> doi: 10.1109/TPAMI.2008.76
- McKinney, W. (2010). Data Structures for Statistical Computing in Python. In S. van der Walt & J. Millman (Eds.), *Proceedings of the 9th python in science conference* (Vol. 1697900, pp. 51–56). Retrieved from <http://conference.scipy.org/proceedings/scipy2010/mckinney.html>

- Nister, D., & Stewenius, H. (2006). Scalable recognition with a vocabulary tree. In *Proceedings of the IEEE computer society conference on computer vision and pattern recognition* (Vol. 2, pp. 2161–2168). IEEE. Retrieved from <http://ieeexplore.ieee.org/document/1641018/> doi: 10.1109/CVPR.2006.264
- Nun, I., Pichara, K., Protopapas, P., & Kim, D.-W. (2014). Supervised Detection of Anomalous Light Curves in Massive Astronomical Catalogs. *The Astrophysical Journal*, 793(2), 16pp. doi: 10.1088/0004-637X/793/2/23
- Nun, I., Protopapas, P., Sim, B., Zhu, M., Dave, R., Castro, N., & Pichara, K. (2015, may). FATS: Feature Analysis for Time Series. , 1–13. Retrieved from <http://arxiv.org/abs/1506.00010>
- Oliphant, T. E. (2007). *Python for scientific computing* (Vol. 9) (No. 3). Retrieved from <http://www.scipy.org/> doi: 10.1109/MCSE.2007.58
- Pichara, K., & Protopapas, P. (2013, nov). Automatic Classification of Variable Stars in Catalogs With Missing Data. *The Astrophysical Journal*, 777(2), 83. Retrieved from <http://stacks.iop.org/0004-637X/777/i=2/a=83?key=crossref.882986e891fcd96a46046022c23135dc> doi: 10.1088/0004-637X/777/2/83
- Pichara, K., Protopapas, P., Kim, D. W., Marquette, J. B., & Tisserand, P. (2012). An improved quasar detection method in EROS-2 and MACHO LMC data sets. *Monthly Notices of the Royal Astronomical Society*, 427(2), 1284–1297. doi: 10.1111/j.1365-2966.2012.22061.x
- Pichara, K., Protopapas, P., & León, D. (2016, feb). Meta-Classification for Variable Stars. *The Astrophysical Journal*, 819(1), 18. Retrieved from <http://stacks.iop.org/0004-637X/819/i=1/a=18?key=crossref.1de9d5e4e8b60e0cd9e40d9ea6802b46> doi: 10.3847/0004-637X/819/1/18
- Popivanov, I., & Miller, R. (2002). Similarity search over time-series data using wavelets. *Proceedings 18th International Conference on Data Engineering*, 212–221. Retrieved from <http://ieeexplore.ieee.org/lpdocs/epic03/wrapper.htm?arnumber=994711> doi: 10.1109/ICDE.2002.994711

Rakthanmanon, T., Campana, B., Mueen, A., Batista, G., Westover, B., Zhu, Q., ... Keogh, E. (2013). Addressing Big Data Time Series: Mining Trillions of Time Series Subsequences Under Dynamic Time Warping. *Transactions on Knowledge Discovery from Data (TKDD)*, 7(3), 3047–3051. Retrieved from [http://portal.acm.org/citation.cfm?id=2513092.2500489](http://portal.acm.org/citation.cfm?id=2513092.2500489&coll=DL&dl=ACM&CFID=254381704&CFTOKEN=61777500&5Cnpapers3://publication/doi/10.1145/2513092.2500489) doi: 10.1145/2500489

Richards, J. W., Starr, D. L., Butler, N. R., Bloom, J. S., Brewer, J. M., Crellin-Quick, A., ... Zechmeister, M., M. & Kürster (2011, may). on Machine-Learned Classification of Variable Stars With Sparse and Noisy Time-Series Data. *The Astrophysical Journal*, 733(1), 10. Retrieved from <http://stacks.iop.org/0004-637X/733/i=1/a=10?key=crossref.b01929a3dbee1845bb80f17663f565fa> doi: 10.1088/0004-637X/733/1/10

Salton, G., & McGill, J. M. (1983). Introduction to modern information retrieval. , 24–51.

Sart, D., Mueen, A., Najjar, W., Keogh, E., & Niennattrakul, V. (2010, dec). Accelerating dynamic time warping subsequence search with GPUs and FPGAs. In *Proceedings - iee international conference on data mining, icdm* (pp. 1001–1006). IEEE. Retrieved from <http://ieeexplore.ieee.org/document/5694075/> doi: 10.1109/ICDM.2010.21

Showcase, R., Cmu, ., Yi, B.-K., Jagadish, H. V., & Faloutsos, C. (1997). Efficient Retrieval of Similar Time Sequences Under Time Warping. *Data Engineering*, 201–208. Retrieved from <http://repository.cmu.edu/compsci>

Sivic, Zisserman, Sivic, J., & Zisserman, A. (2003). Video Google: A text retrieval approach to object matching in videos. *Toward Category-Level Object Recognition*, 1470–1477. Retrieved from <http://www.robots.ox.ac.uk/~vgg/publications/papers/sivic03.pdf><http://www.robots.ox.ac.uk/~vgg/publications/papers/sivic03.pdf><http://www.cse.unr.edu/~bebis/CS773C/ObjectRecognition/Papers/VideoGoogle.pdf><http://ieeexplore.ieee.org/document/1238663/> doi: 10.1109/ICCV.2003.1238663

- Szalkai, B. (2013, apr). *An implementation of the relational k-means algorithm*. Retrieved 2017-04-30, from <http://arxiv.org/abs/1304.6899><http://arxiv.org/pdf/1304.6899v1.pdf>
- Udalski, A., Szymański, M., Kałużny, J., Kubiak, M., & Mateo, M. (1996). The optical gravitational lensing experiment. *Acta Astronomica*, 42(4), 253–284.
- Udalski, A., Szymański, M. K., Soszyński, I., & Poleski, R. (2008, jul). The optical gravitational lensing experiment. Final reductions of the OGLE-III data. *Acta Astronomica*, 58(2), 69–87. Retrieved from <http://arxiv.org/abs/0807.3884>
- Van Der Walt, S., Colbert, S. C., & Varoquaux, G. (2011, mar). The NumPy array: A structure for efficient numerical computation. *Computing in Science and Engineering*, 13(2), 22–30. Retrieved from <http://ieeexplore.ieee.org/document/5725236/> doi: 10.1109/MCSE.2011.37
- Wang, X., Mueen, A., Ding, H., Trajcevski, G., Scheuermann, P., & Keogh, E. (2013). Experimental comparison of representation methods and distance measures for time series data. *Data Mining and Knowledge Discovery*, 26(2), 275–309. doi: 10.1007/s10618-012-0250-5
- Wu, Y.-L., Agrawal, D., & El Abbadi, A. (2000). A Comparison of DFT and DWT Based Similarity Search in Time-series Databases. In *Proceedings of the 9th international conference on information and knowledge management (cikm)* (Vol. 35, pp. 488–495). New York, New York, USA: ACM Press. Retrieved from <http://portal.acm.org/citation.cfm?doid=354756.354857> doi: 10.1145/354756.354857