



HAL
open science

Explainable classification of astronomical uncertain time series

Michael Franklin Mbouopda, Emille E O Ishida, Engelbert Mephu Nguifo,
Emmanuel Gangler

► **To cite this version:**

Michael Franklin Mbouopda, Emille E O Ishida, Engelbert Mephu Nguifo, Emmanuel Gangler. Explainable classification of astronomical uncertain time series. 2022. hal-03790123

HAL Id: hal-03790123

<https://uca.hal.science/hal-03790123>

Preprint submitted on 28 Sep 2022

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

EXPLAINABLE CLASSIFICATION OF ASTRONOMICAL UNCERTAIN TIME SERIES

Michael Franklin MBOUOPDA

LIMOS
University Clermont Auvergne
Clermont-Ferrand
michael.mbouopda@uca.fr

Emille E. O. ISHIDA

Laboratory of Physics of Clermont
University Clermont Auvergne
Clermont-Ferrand
emille.ishida@clermont.in2p3.fr

Engelbert MEPHU NGUIFO

LIMOS
University Clermont Auvergne
Clermont-Ferrand
engelbert.mephu_nguifo@uca.fr

Emmanuel GANGLER

Laboratory of Physics of Clermont
University Clermont Auvergne
Clermont-Ferrand
emmanuel.gangler@clermont.in2p3.fr

ABSTRACT

Exploring the expansion history of the universe, understanding its evolutionary stages, and predicting its future evolution are important goals in astrophysics. Today, machine learning tools are used to help achieving these goals by analyzing transient sources, which are modeled as uncertain time series. Although *black-box* methods achieve appreciable performance, existing interpretable time series methods failed to obtain acceptable performance for this type of data. Furthermore, data uncertainty is rarely taken into account in these methods. In this work, we propose an uncertainty-aware subsequence based model which achieves a classification comparable to that of state-of-the-art methods. Unlike conformal learning which estimates model uncertainty on predictions, our method takes data uncertainty as additional input. Moreover, our approach is explainable-by-design, giving domain experts the ability to inspect the model and explain its predictions. The explainability of the proposed method has also the potential to inspire new developments in theoretical astrophysics modeling by suggesting important subsequences which depict details of light curve shapes. The dataset, the source code of our experiment, and the results are made available on a public repository.

Keywords Time series · Classification · Explainability · Astronomy · Uncertainty

1 Introduction

Machine learning (ML) has become an ineluctable tool for analyzing and extracting meaningful information from data. Classically exclusively applied on tabular data, it is nowadays also effective on image, video, text, and also time series data. The latter is the type of data we will focus on in this paper. Specifically, this work is about *time series classification*, a ML task whose goal is to learn a function (i.e a classifier) that maps time series to a set of discrete classes. A time series is an ordered and finite sequence of values. Some examples of time series are the daily COVID cases and the monthly groundwater level. Time series classification has been applied in several domains including online harassment detection [1], medicine [2, 3], emotion recognition [4], anomaly detection [5], and in physics [6, 7, 8, 9]. This usability is facilitated by toolkits such as Sktime [10], which unifies the existing time series classification algorithms under the same user-friendly API. However, the existing methods are generally not applicable to *uncertain* time series; In fact, as far as we know, the uncertain shapelet transform (or simply UST) method [11] is the only one that has been designed for uncertain time series classification.

An uncertain time series (uTS) is a time series of *imprecise* values. Unlike a *regular* time series which is an ordered sequence of real numbers, an uTS is a sequence of pairs of numbers such that the first number of a pair is the best




Figure 1: Uncertain time series illustrations

estimate and the second one is the error on that estimate; therefore the exact values of an uTS are unknown. Figure 1a illustrates a simulated uTS: the blue line is the best estimate and the vertical red bars represent the uncertainty intervals (i.e the exact unknown values are somewhere on the vertical red bars). Figure 1b is a real uTS extracted from the PLAsTiCC [6] dataset: any time series that lies in the red region could be the exact unknown time series. uTS classification should not be confused with an application of conformal learning; In fact, in conformal learning, the input data is assumed to be uncertainty-free, and the goal is to compute the uncertainty of model’s predictions. On the opposite, the goal of uTS classification is to infer a classifier from and for uncertain data.

Uncertain time series are preponderant in transient astrophysics. Astronomical objects whose brightness vary with time (a.k.a transients) are primarily characterized by the presence or absence of specific chemical elements found in their spectra. This data taking process (called spectroscopy) is very time-consuming and requires very good observation conditions to be performed. Moreover, since transients are objects which appear in the sky for a limited period of time then disappear forever, there is a small time-window of opportunities when such measurements can be taken.

Alternatively, we can also associate different classes of astronomical transients to the respective shape of their light curves (brightness variation as a function of time). In this case, we need to repeatedly measure the brightness of the source in a relatively broad region of the wavelength spectrum. This process, called photometry, is less expensive and imposes more manageable constraints on observation conditions. However, measurements are more prone to uncertainties (due to moonlight, twilight, clouds, etc) in the flux determination and the distinction between light curves from different classes is subtle, resulting in less accurate classifications. Nevertheless, since there is not enough spectroscopic resources to provide definite label for all photometric observed objects, being able to effectively analyze uncertain photometric light curves means that a wider range of the universe can be quickly understood and at a lower cost.

The Vera C. Rubin Observatory¹ is a ground-based observatory, currently under construction in Chile, whose goal is to conduct the 10-year Legacy Survey of Space and Time (LSST) in order to produce the deepest and widest images of the universe. The observatory is expected to start producing data in early 2024, and in order to prepare the community for the arrival of its data, one important data challenge was put in place: the Photometric LSST Astronomical Time-Series Classification Challenge or simply PLAsTiCC [6]. The goal was to identify machine learning models able to classify 14 types of transients in simulated data, represented by uncertain time series, or light curves. The ultimate goal behind the challenge was to understand which methods are expected to perform better in LSST-like data, thus preparing the community to the arrival of its data and help understanding the universe’s expansion history. Therefore, using interpretable approaches was very important. However, contributors focused on minimizing the classification loss by employing techniques such as mixture of classifiers and data augmentation [12] while neglecting explainability. In this paper, we address this problem with explainability in mind.

We consider two approaches to classify uTS in an explainable manner : the first one ignores uncertainty and uses only the best estimates, while the second one takes uncertainty into account. Ignoring uncertainty makes the task a *regular* time series classification task, allowing the usage of Shapelet Transform Classification or simply STC [13], an effective and explainable *regular* time series classification algorithm. This model failed to find any valid shapelet on PLAsTiCC, and therefore could not perform the classification task. We performed extensive hyper-parameter tuning tests, but the

¹<https://lsst.org/>

result was the same. We also tried to take uncertainty into account by using the Uncertain Shapelet Transform algorithm [11], but as expected, this method also failed since it is an extension of STC for uncertain time series.

In this paper, we propose the Uncertain Scalable and Accurate Subsequence Transform (or uSAST for short) method which is able to achieve an F1-score of 70% while providing faithful explanation similarly to STC. The rest of this paper is organized as follows: we start by presenting the background and the related works. We continue by describing the uSAST method. Finally, we detail our experiments and the obtained results before concluding this work.

2 Background

Definition 1 (Time series). *A time series (TS) of length m is a finite sequence of ordered values.*

$$T = (t_1, t_2, \dots, t_m), t_i \in \mathbb{R}, m > 0$$

Definition 2 (Uncertain time series). *An uncertain time series (uTS) is defined similarly to a time series, but each value has an uncertainty represented by a positive real number.*

$$T = (t_1 \pm \delta t_1, t_2 \pm \delta t_2, \dots, t_m \pm \delta t_m), t_i \in \mathbb{R}, m > 0, \delta t_i \in \mathbb{R}_+$$

Definition 3 (Subsequence). *A subsequence (respectively an uncertain subsequence) is a sequence of consecutive values extracted from a TS (respectively an uTS).*

Definition 4 (Distance). *The distance between a subsequence S of length l and a time series of length m is defined as follows:*

$$Dist(S, T) = \min_{P \in T^l} dist(S, P),$$

where $T^l = \{(t_i, t_{i+1}, \dots, t_{i+l}) \mid 1 \leq i \leq m - l + 1\}$

The $dist(\cdot, \cdot)$ function in Definition 4 could be any distance metric. In practice the Euclidean Distance (ED) and the Dynamic Time Warping (DTW) are generally used. The definition is also applicable between uTS and uncertain subsequence by ignoring the uncertainty or by taking it into account using an uncertain distance, the UED distance [11].

Definition 5 (Uncertain Euclidean Distance). *The Uncertain Euclidean Distance (UED) between two uncertain subsequences S_1 and S_2 of same length l is defined as:*

$$UED(S_1, S_2) = \sum_{i=1}^l (s_{1,i} - s_{2,i})^2 \pm 2 \sum_{i=1}^l |s_{1,i} - s_{2,i}| (\delta s_{1,i} + \delta s_{2,i})$$

Let $D = \{(T_i, c_i) \mid 1 \leq i \leq n\}$ be a dataset of n time series T_i (respectively uncertain time series) with their class labels c_i taken from a discrete finite set C such that the cardinality of C is much less than n . We can define the notions of *separator* and *shapelet* for this dataset.

Definition 6 (Separator). *A separator (respectively uncertain separator) is a pair of a subsequence S (respectively uncertain subsequence) and a threshold ϵ that divide the dataset in two groups D_{left} and D_{right} such that:*

$$D_{left} = \{(T_i, c_i) \mid Dist(S, T_i) < \epsilon, 1 \leq i \leq n\}$$

$$D_{right} = \{(T_i, c_i) \mid Dist(S, T_i) \geq \epsilon, 1 \leq i \leq n\}$$

Definition 7 (Shapelet). *A shapelet (respectively uncertain shapelet) is a separator (respectively uncertain separator) that maximizes the information gain similarly to splitting nodes in decision trees [14].*

3 Related works

Time series classification is performed regarding *global* features, *local* features, or both. Historically, only global features were considered; in particular, the classification was done using the one nearest neighbor (1-NN) classifier and the Dynamic Time Warping (DTW) distance. The Elastic Ensemble (EE) is an improvement of the *global* features classification, obtained by ensembling several distance measures [15]. The Fast Ensemble of Elastic Distances (FastEE) significantly reduces the computation time of the Elastic Ensemble [16].

Local feature-based methods are organized as dictionary-based, interval-based or subsequence-based. Dictionary-based methods proceeds by representing each time series using a finite set of discrete symbols using techniques such as Symbolic Fourier Approximation (SFA) [17] and Symbolic Aggregate approXimation (SAX) [18]. Some methods that implement these techniques are BOSS [19], MUSE [20] and TDE [21]. Interval-based methods assume that the

whole time series is not relevant for classification, but that only some segments (i.e intervals) contain the discriminative features. The first step in these methods is the identification of the relevant intervals, then a set of features (mean, median, slope, ...) are computed for each interval and finally a supervised classifier is trained on the computed features. Some methods that use this approach are TSF [22], CIF [23] and STSF [24]. Subsequence-based methods also assume that only some segments are relevant for classification, but unlike interval-based methods which use phase-dependent features (i.e. the locations of the intervals are fixed), subsequence-based methods use phase-independent features. Subsequence-based methods perform in three steps: first, the relevant subsequences are identified, then each time series is transformed to a vector of its distances to the relevant subsequences, and finally a supervised classifier is trained on the obtained vectors. Some of these methods are Shapelet-based decision trees [14], Shapelet Transform [13], Uncertain Shapelet Transform [11] and SAST [25]. Uncertain Shapelet Transform (UST) is, to our knowledge, the only subsequence-based method that supports uTS classification.

The most accurate methods for time series classification are TS-CHIEF [26], which combines both local and global features in a tree-based classifier, HIVE-COTE [15, 27], which combines several methods to extract local and global features, and ROCKET [28, 29], which employs random convolutional kernels to extract local features. These methods are also known to have explainability issues. Subsequence-based methods are the most easier to explain; However, all of them failed on the PLAsTiCC dataset. In the next section, we will describe an accurate and interpretable method which produces competitive results on this data set and can also be easily applied to any other uTS dataset.

4 Uncertain Subsequence Transform Classification

In this section, we describe a new uncertain time series classification method based on uncertainty propagation as in UST [11] and subsequence transform as in SAST [25]. In fact, uncertainty propagation is an effective approach to analyze uncertain data [30, 31] and particularly uncertain time series [11]. By using a single random instance from each class, SAST is more scalable and at least as accurate as STC [25] while keeping STC interpretability capabilities.

Given a time series dataset, SAST follows four steps: *i*), one instance is randomly selected from each class: these are called reference time series; *ii*) a set containing every subsequences from the selected time series is created; *iii*) each instance in the dataset is replaced by the vector of its distances to each subsequence obtained in the second step; *iv*) a supervised classifier is trained on the transform dataset.

Performing classification following the SAST steps could be inefficient because of the redundancy in the set of subsequences obtained at the second step. The redundancy is particularly high for small length subsequences and in datasets such as electrocardiogram (ECG) and PLAsTiCC, in which repetitive patterns occur very often. Furthermore, the third step is based on the application of Definition 4 using the Euclidean distance and, therefore, only the most similar subsequence is considered; however, taking into account the number of occurrences of the best match is important in some contexts. To overcome these limitations, we define the notion of ϵ -similarity as follows:

Definition 8 (ϵ -similarity). *Two subsequences (respectively uncertain subsequences) S_1 and S_2 of same length l are ϵ -similar if the distance between them is less than or equal to a user-defined threshold $\epsilon \geq 0$.*

$$\epsilon\text{-similar}(S_1, S_2) = \begin{cases} True, & \text{if } dist(S_1, S_2) \leq \epsilon \\ False, & \text{otherwise} \end{cases}$$

Theorem 1. *The ϵ -similar relationship is not transitive.*

Proof. Let X , Y , and Z be three subsequences of same length l such that $\epsilon\text{-similar}(X, Y) = True$ and $\epsilon\text{-similar}(Y, Z) = True$. Let us assume that the transitivity property is verified, that is $\epsilon\text{-similar}(X, Z) = True$. A counterexample is built by considering X , Y , and Z as points in a high dimensional space (\mathbb{R}^l) such that $dist(X, Y) = dist(Y, Z) = \epsilon$, and $XY \perp XZ$. The following derivation proves the theorem:

$$\begin{aligned} dist(X, Z) &= \sqrt{dist(X, Y)^2 + dist(Y, Z)^2} \\ &= \sqrt{\epsilon^2 + \epsilon^2} \\ &= \epsilon\sqrt{2} \\ &> \epsilon \\ &\implies \epsilon\text{-similar}(X, Z) = False \end{aligned}$$

□

Using Definition 8, we can reduce redundancies and count subsequence frequencies in SAST. The updated SAST method, hereafter SAST+, is detailed in Algorithm 1.

Algorithm 1 SAST+

Require: $D = \{(T_1, c_1), (T_2, c_2), \dots, (T_n, c_n)\}$, k : the number of instances to use per class, $length_list$: the list of subsequence lengths, C : the classifier to use, ϵ : ϵ -similarity parameter.

- 1: $D_c \leftarrow randomlySelectInstancesPerClass(D, k)$ ▷ Randomly select k instances per class from the dataset
▷ Generate every patterns of length in $length_list$ from D_c , using ϵ to remove similar patterns
- 2: $S \leftarrow generateSubsequences(D_c, length_list, \epsilon)$
- 3: $D_f \leftarrow \emptyset$
- 4: **for** $i \leftarrow 1$ **to** n **do** ▷ Transformed the dataset using every patterns in S
- 5: $x_i \leftarrow []$
- 6: **for** $j \leftarrow 1$ **to** $|S|$ **do** ▷ The procedure $distAndCount(T_i, S_j, \epsilon)$ returns $Dist(T_i, S_j)$ and the number of occurrences of the subsequence S_j in T_i
▷ The procedure $distAndCount(T_i, S_j, \epsilon)$ returns $Dist(T_i, S_j)$ and the number of occurrences of the subsequence S_j in T_i
- 7: $x_i[j], x_i[j + |S|] \leftarrow distAndCount(T_i, S_j, \epsilon)$
- 8: **end for**
- 9: $D_f \leftarrow D_f \cup \{(x_i, c_i)\}$
- 10: **end for**
- 11: $clf \leftarrow trainClassifier(C, D_f)$ ▷ Train the classifier on the transformed dataset
- 12: **return** (clf, S) ▷ The trained classifier and the subsequences

The time complexity of the SAST method is $O(N_c) + O(kN_cm^2) + O(nm^3) + O(classifier)$, where N_c is the number of classes, n the number of time series, m the length of the time series and k the number of reference time series per class [25]. In practice, it is not necessary to have k greater than one. Removing redundancies in SAST is done only once (during the training phase) with a theoretical time complexity of $O(km^4)$; counting frequencies is done while computing the distance in a constant time. Therefore, the SAST+ time complexity is $O(N_c) + O(kN_cm^2) + O(nm^3) + O(classifier) + O(km^4)$ which is asymptotically equivalent to $O(classifier) + O(km^4)$. Removing redundancies makes SAST+ much faster than SAST during inference.

Similarly to the Uncertain Shapelet Transform [11], the uncertain SAST+ (uSAST+) is obtained by using UED as the distance metric in Algorithm 1; allowing uncertainties to be propagated to the classifier which then uses these uncertainties to learn robust decision boundaries.

5 Experiment

5.1 The PLAsTiCC dataset

As far as we know, existing methods published on uTS classification have never been evaluated on real uncertain time series datasets, but solely on simulated datasets. The corresponding simulated datasets have never been made publicly accessible neither for reproducibility reasons, nor for facilitating research on uTS. In this work, we evaluate our method on a realistic publicly available uncertain time series dataset from the astrophysics domain.

The Photometric LSST Astronomical Time-Series Classification Challenge (PLAsTiCC) dataset contains uncertain time series representing the brightness evolution of astronomical transients including supernovae, kilonovae, active galactic nuclei and eclipsing binary systems [6], among others. Each object is represented as a multivariate uncertain time series of 6 dimensions named u, g, r, i, z, y , each corresponding to a particular broadband wavelength filter. After the challenge was finished, the organizers made available an updated version of the data through Zenodo² with some bug fixes and the classification answers for both the training and test sets. In this work, we demonstrate our method using only uncertain time series from the training set, but the methodology is general enough to be extended to the test set. There are 7848 transients in the dataset, grouped in 14 different classes, and the number of objects in the classes are highly imbalanced. More specifically, the most underpopulated class has only 0.3% of objects, whereas the most populated one contains 29% of the objects. Furthermore, the dataset contains a lot of missing observations. We handled this with the help of astrophysicists who suggested to fill missing data using a rolling average with a window of length 5. Missing values and corresponding error bars are replaced by the mean and standard deviation of the window. This procedure translated the original dataset into a homogeneously sampled uncertain time series. The preprocessed dataset is made public³.

²<https://zenodo.org/record/2539456>

³Cleaned dataset: <https://drive.uca.fr/f/f0741be3fb77402f8e82/>

Our implementation uses the Python programming language and is based on the Scikit-learn machine learning library [32] and the Sktime time series dedicated machine learning library [10]. The experiment is run on a computing node equipped with 1 Gb of RAM and an AMD EPIC 7452 processor containing 64 logical cores of 2.35 GHz frequency. The source code of our experiments and all the results we discuss in this paper are publicly available on GitHub⁴.

5.2 Results

Since PLAsTiCC is a multivariate uncertain time series dataset, the subsequence transformation is performed on each dimension independently. The transformations from each dimension are then concatenated together to build a large matrix which is subsequently fed to the supervised classifier. We used 80% of the data for training and the remaining is used for testing.

5.2.1 Shapelet-based methods results:

Shapelet-based classification is a special case of subsequence-based classification which consider only shapelets as relevant subsequences. We considered two shapelet-based methods STC [13] and UST [11] for their interpretability. For both methods, we kept every parameters to their default values except the *minimum information gain* parameter which is the threshold used to decide if a separator is a valid shapelet. We tried different values for this parameter without success, none of these methods were able to find a single valid shapelet in the dataset. Since feature extraction was not successful, classification was not possible. This result is due to the dataset being highly imbalanced and the uncertain time series from different classes being too similar in shape. The same dimension of two randomly selected samples from two different classes is shown on Figure 2. The left figure which is a Supernova Type Ia-x (SNIax) looks like a left-shifted version of the right figure which is a Supernova Type Ia-91bg (SNIa-91bg). SNIax and SNIa-91bg are known to be difficult to distinguish by astrophysicists. This observation holds, with different magnitude, for other classes in the PLAsTiCC dataset and therefore, any shapelet-based methods might struggle to find shapelets in this dataset.




Figure 2: Two supernova from PLAsTiCC. They look similar in terms of shapes although they are from distinct classes.

5.2.2 SAST-based methods results:

For this experiment we considered different SAST+ configurations in order to measure the effect of taking uncertainty into account, dropping duplicates and counting the number of occurrences of patterns (i.e patterns frequency). We named configurations that ignore uncertainty as SAST<X> and those which take uncertainty into account as uSAST<X>, where <X> is either : *i*) an empty string to specify that duplicate subsequences are not removed and the patterns frequency is ignored; *ii*) the character *d*, meaning that duplicate patterns are removed; *iii*) the string *dc*, meaning that duplicate patterns are removed and the frequency of patterns is taken into account.

We use three different supervised classifiers, namely Random Forest (RF), eXtreme Gradient Boosting (XGBoost) and the Ridge regression with Leave-One-Out cross-validation (RidgeCV). The cross-validation procedure is used to find the best regularization parameter. We set the minimum and maximum subsequence lengths to 20 and 60 respectively, with a step of 10. Compared to a step of 1, a step of 10 reduces the chance of having similar subsequences while reducing the number of subsequences to be used. We observed that the classification performance is better with this

⁴Source code: <https://anonymous.4open.science/r/usast-FBC0/>

setup as can be seen in the supplementary material. The ϵ -similarity is computed with $\epsilon = 0.25$ as experiments shown that too much relevant subsequences are discarded with higher values. The parameters of the classifiers are left to their default values, except for the regularization parameter in RidgeCV which is selected using cross-validation. As the reference time series are chosen randomly, we run each experiment 3 times and we report the average precision, recall, F1 score, cross entropy loss and the time taken for training and inference (in hours). As PLAsTiCC is an imbalanced multiclass dataset, we use a weighted average to compute the precision, recall and F1 score; the weights being the percentage of each class in the dataset. Table 1 shows the result using the XGBoost classifier only as it has led to the best classification performance. However, detailed results are available in the supplementary material.

Table 1: Results on PLAsTiCC averaged over 3 runs.

	Precision	Recall	F1 score	LogLoss	Time (h)
uSAST	0.72 ± 0.01	0.72 ± 0.00	0.69 ± 0.01	0.96 ± 0.01	51.03 ± 0.12
uSASTd	0.72 ± 0.00	0.73 ± 0.00	0.70 ± 0.01	0.97 ± 0.01	43.49 ± 0.27
uSASTdc	0.71 ± 0.01	0.72 ± 0.01	0.69 ± 0.01	0.96 ± 0.01	43.52 ± 0.72

The first observation is that any variant of our proposed method is able to achieve around 70% precision, recall and F1 score, unlike shapelet-based methods which completely failed on the PLAsTiCC dataset. This result corroborates with the claim that pruning subsequences before the effective classification could sometimes lead to poor performance [25]. Dropping duplicates, counting patterns frequency or doing both does not have significant impact on the classification performance. However, dropping duplicate makes the models faster. In particular, uSASTd is about 12 hours faster than uSAST. Counting pattern frequency does not add a computation overhead because it is done while computing the distance in $O(1)$ time.

Choosing the right subsequence lengths to considered is challenging and assessing all possible values is computationally expensive; However, domain knowledge could guide in setting this parameter as it is application-dependent.

PLAsTiCC contains objects that are either galactic or extra-galactic, and whose light curves were obtained following a Deep Drilling Fields (DDF) or Wide Fast Deep (WFD) observation strategy. Extra-galactic objects are further away than galactic ones, they are fainter and more difficult to be observed. DDF light curves contain more frequent observation points than WFD ones. Thus, DDF light curves provide a more certain determination of the time series properties than their WFD counterparts which have more uncertainties. Table 2 gives the performances of the model uSASTd regarding if the objects are galactic or not, DDF or WFD. The model is considerably better at classifying galactic objects than extra-galactic ones, and a little better at classifying DDF objects than WFD ones. While the model achieves an F1 score of 94% for galactic objects in DDF, it achieves an F1 score of only 67% for extra-galactic objects in WFD. This is directly related to the astrophysical nature of galactic objects. These are, in general, variables whose brightness go through many cycles within the 3 years covered by our data. On the other hand, extragalactic objects are dominated by transients, consisting of only 1 region of signal which never repeats, thus rendering a smaller quantity of information encoded in its time series.

Table 2: uSASTd performance regarding if the object are galactic or extra-galactic, are from the DDF or WFD.

		Galactic	Extra-galactic	Both
DDF	Precision	0.96	0.73	0.77
	Recall	0.94	0.75	0.79
	F1 score	0.94	0.71	0.76
WDF	Precision	0.94	0.67	0.71
	Recall	0.84	0.64	0.71
	F1 score	0.87	0.61	0.67
Both	Precision	0.94	0.68	0.72
	Recall	0.86	0.67	0.73
	F1 score	0.88	0.64	0.70

The data set includes 6 classes with overall similar behavior (42, 52, 62, 67, 90, 95). Among these, astronomers are specially interested in type 90 (SNIa), which is used as distance indicator in cosmological analysis [33]. Reporting our results as a binary problem with class 90 against all others, we achieve 85% precision, 81% recall and 82% F1 score. Therefore, our method is able to correctly classify a high proportion of SNIa despite its similar behavior to other classes.

5.2.3 Ablation study:

Here, we study the impact of taking uncertainty into account. In particular, we compare the results obtained when uncertainty is ignored (Table 3) to the results obtained when uncertainty is taken into account (Table 1).

Table 3: Results on PLAsTiCC averaged over 3 runs when uncertainty is ignored.

	Precision	Recall	F1 score	LogLoss	Time (h)
SAST	0.65 ± 0.01	0.67 ± 0.00	0.63 ± 0.00	1.16 ± 0.01	16.41 ± 0.52
SASTd	0.66 ± 0.02	0.68 ± 0.00	0.64 ± 0.00	1.14 ± 0.00	12.79 ± 0.84
SASTdc	0.66 ± 0.01	0.68 ± 0.00	0.64 ± 0.01	1.14 ± 0.01	12.99 ± 0.30

Taking uncertainty into account increases the classification performance in terms of precision, recall, F1 score and cross entropy loss. In fact, from SASTd to uSASTd, there is a gain of 6% in precision, 5% in recall, 6% in F1 score. It can also be seen that the model is more confident on its predictions as the loss has decreased. However, this gain in performance requires almost four times more computation.

5.2.4 Comparison to SOTA:

In this subsection, we compare our proposed method to the state-of-the-art multivariate time series classification methods ROCKET [28], MUSE [20] and XEM [34] which have been shown to be among the most accurate methods for this task [35]. Results are shown in Table 4.

Table 4: uSASTd vs SOTA results.

	Precision	Recall	F1 score	Time (h)
uSASTd	0.72 ± 0.00	0.73 ± 0.00	0.70 ± 0.01	43.49 ± 0.27
MUSE	0.71 ± 0.01	0.73 ± 0.01	0.71 ± 0.01	3.36 ± 0.04
ROCKET	0.77 ± 0.00	0.77 ± 0.00	0.75 ± 0.00	0.05 ± 0.00
XEM	0.69 ± 0.01	0.71 ± 0.00	0.69 ± 0.00	12.24 ± 0.46

The classification performance of our method is comparable to those of the SOTA methods. In particular, uSASTd achieves better precision, recall and F1 score compared to XEM on PLAsTiCC. uSASTd and MUSE have similar classification performance. ROCKET achieves the best classification performance. SOTA methods are faster than our proposal. Except for XEM which is explainable-by-design, SOTA methods are not explainable. In fact, ROCKET uses the proportion of positive values obtained after applying random convolutions. MUSE uses bag of words obtained after applying some transformations to the time series. These features have no particular meaning for domain experts. Our method does not have this limitation, as it is based on features that are intelligible to domain experts.

5.2.5 Explainability:

One of the best properties of subsequence-based classification is its interpretability. The explanation could be done either locally, when it concerns only a single instance, or globally when it concerns the whole model. In any case, this is generally done by inspecting the model in order to extract the most discriminative subsequences [14]. These subsequences could also be found using a post-hoc method such as LIME [36] or SHAP [37], but since our approach is explainable-by-design, inspecting the model is sufficient. More specifically, since the classifier used in our model is tree-based, the information gain can be used as a measure of the discriminative power of the subsequences similarly to what is done in shapelet-based methods. The local explainability of our method is obtained by inspecting the subsequence on which the model focused the most in order to make the prediction for a single instance. Figure 3 shows local explanations for a Supernova Type Ia (SNIa) and a Core-collapse Supernova Type II-P (SNII-P) correctly classified by the model. The ‘‘P’’ in the denomination of the latter references the plateau phase observed in its time-series just after maximum brightness. This feature is clearly shown in the bottom panel of Figure (3). This confirms that our model focuses on the relevant regions and dimensions of the time series to make the classification. Being able to correctly learn the dimension’s relevance is crucial as the discriminative subsequence may appear only in a subset of the dimensions. Furthermore, the location of the discriminative subsequence may not be the same on every dimension. In PLAsTiCC in fact, depending how far is the object, the light may be visible only on some wavelengths (i.e. dimension).

Due to the accelerated expansion of the universe, objects which are further away are also moving with a higher velocity. Thus, there is a Doppler effect in the observed light which shifts it to higher wavelengths. Thus, closer (galactic) objects will generally have higher signals in lower wavelengths than further away (extragalactic) ones. Our method perfectly captures the Doppler effect unlike XEM which cannot identify from which dimensions the discriminative subsequences is located.

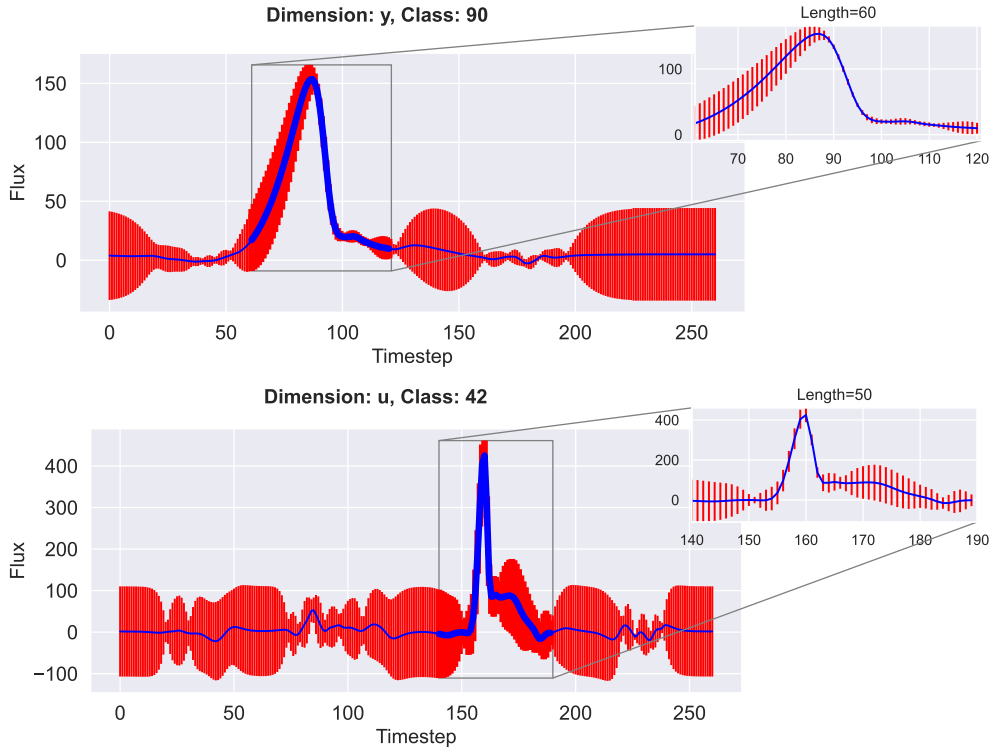


Figure 3: Local explainability of a Supernova Type Ia (top) and a Core-collapse Supernova Type II (bottom)

A global explanation is obtained by building a subsequence-based profile of each of the class. The top 20 most discriminative subsequences from the uSASTd model are shown in Figure 4. Subsequences that are from the same class label are plotted with the same color, its rank, its class label and its type are given at the top of its corresponding plot. The type is either *Value* if the discriminative power comes from the value itself or *Uncertainty* if the discriminative powers comes from the uncertainty. The dimension from which the subsequences are coming from are also given on the figure.

It is observed that the discriminative power is generally due to the value, but sometimes it is due to the uncertainty (for example subsequences #18 and #19). Seeing that some subsequences are important because of their uncertainty emphasizes the fact that taking uncertainty into account is important and improves the classification performance. There are also some subsequences that are too similar despite the fact that duplicate subsequences have been dropped; for instance, the subsequences #3 and #7. This is because the similarity between subsequences is computed using the Uncertain Euclidean Distance (UED) which considers the subsequences to be perfectly aligned. This problem can be resolved by using an elastic distance such as the DTW distance at the cost of more computational time since such distances generally have at least quadratic time complexity while UED is linear. From the domain knowledge point of view, these discriminative subsequences are able to grasp the important shapes commonly associated with their respective class of astronomical transients. Subsequences #1 and #6 were taken from class 16 (eclipsing binary) and clearly show the expected light curve from a well measured binary system where one star eclipses the other exactly in the line of sight, thus leading to a decrease in brightness. Subsequence #19 is also associated to the eclipsing binary class, but in this case the signal is less clear, corresponding to an object which is further away – thus leading to low signal and large uncertainties. We also call attention to the supernova-like behavior exhibited by subsequences #4 and #9 – one single burst events whose brightness are only visible for weeks to months. The fact that such characteristic behaviors are easily spotted in the list of most important subsequences certifies that our final classification results are in line with the expert definition of such classes and hence, shows that our model is safe and trustworthy. Moreover,




Figure 4: The top 10 most discriminative subsequences in the PLAsTiCC dataset

further investigations of a more extensive list of important subsequences have the potential to reveal unexpected time series shapes and promote the development of more detail theoretical models for such astrophysical sources.

6 Conclusion and future directions

The classification of time series with available uncertainty measures is an under-explored and challenging task. In this work, we proposed an approach to perform this task with a global F1 score of 70%, without using techniques such as data augmentation nor oversampling. The explainability of the proposed approach allows domain experts to not only understand individual predictions, but also to characterize each class by a set of subsequences with high discriminative power, which can then be used to perform other important tasks in astrophysics such as novel astronomical transients detection and anomaly detection. The ablation study shown the positive impact of taking uncertainty into account. A limitation of the approach is the time complexity, which could be considerably high for datasets with relatively long uncertain time series. A future direction would consist of further reducing the number of subsequences to be used and optimizing the computation time of the method. Another future direction would consist of finding a better way of managing uncertainty during the classification step in order to improve the performances. Nevertheless, the results presented in this work illustrate how our approach is effective in identifying meaningful subsequences which, beyond the classification performance, can provide important information to the expert. The approach is flexible enough to be applied to other scientific domains where uncertain time series are the common, thus enabling future advances in multiple subject areas.

Acknowledgments

This work is funded by the French Ministry of Higher Education, Research and Innovation. Thanks to the TransiXplore project team which helped us understanding and preprocessing the PLAsTiCC dataset and thanks to the Large Synoptic Survey Telescope (LSST) project which published the dataset.

References

- [1] Piotr Janiszewski, Mateusz Lango, and Jerzy Stefanowski. Time aspect in making an actionable prediction of a conversation breakdown. In *ECML/PKDD*, pages 351–364, 2021.
- [2] Elishiah Miller, Zane MacFarlane, Seth Martin, Nilanjan Banerjee, and Ting Zhu. Radar-based monitoring system for medication tampering using data augmentation and multivariate time series classification. *Smart Health*, 23, 2022.

- [3] Qingxiong Tan, Mang Ye, Baoyao Yang, Siqi Liu, Andy Jinhua Ma, Terry Cheuk-Fung Yip, Grace Lai-Hung Wong, and PongChi Yuen. Data-gru: Dual-attention time-aware gated recurrent unit for irregular multivariate time series. In *AAAI*, volume 34, pages 930–937, 2020.
- [4] Sidra Rafique, Nadia Kanwal, Mohammad Samar Ansari, Mamoona Asghar, and Zuhair Akhtar. Deep learning based emotion classification with temporal pupillometry sequences. In *ICECET*, pages 1–6. IEEE, 2021.
- [5] Lifeng Shen, Zhongzhong Yu, Qianli Ma, and James T Kwok. Time series anomaly detection with multiresolution ensemble decoding. In *AAAI*, volume 35, pages 9567–9575, 2021.
- [6] Tarek Allam Jr et al. The photometric lsst astronomical time-series classification challenge (PLAsTiCC): Data set. *arXiv:1810.00001*, 2018.
- [7] Kyle Boone. Avocado: Photometric classification of astronomical transients with gaussian process augmentation. *The Astronomical Journal*, 158(6):257, 2019.
- [8] Anais Möller and Thibault de Boissière. Supernova: an open-source framework for bayesian, neural network-based supernova classification. *MNRAS*, 491(3):4277–4293, 2020.
- [9] Marco Leoni, Emille E. O. Ishida, Julien Peloton, and Anais Möller. Fink: early supernovae Ia classification using active learning. *arXiv:2111.11438v1*, 2021.
- [10] Markus Löning, Anthony Bagnall, Sajaysurya Ganesh, Viktor Kazakov, Jason Lines, and Franz J Király. Sktime: A Unified Interface for Machine Learning with Time Series. In *NeurIPS Workshop*, 2019.
- [11] Michael Franklin Mbouopda and Engelbert Mephu Nguifo. Uncertain time series classification with shapelet transform. In *ICDM Workshop*, pages 259–266. IEEE, 2020.
- [12] R. Hložek et al. Results of the photometric lsst astronomical time-series classification challenge (plasticc). *arXiv:2012.12392v1*, 2020.
- [13] Jon Hills, Jason Lines, Edgaras Baranauskas, James Mapp, and Anthony Bagnall. Classification of time series by shapelet transformation. *DMKD*, 28:851–881, 2014.
- [14] Lexiang Ye and Eamonn Keogh. Time series shapelets. In *SIGKDD*, page 947. ACM, 2009.
- [15] Jason Lines, Sarah Taylor, Anthony Bagnall, and A Bagnall. Time series classification with HIVE-COTE: The hierarchical vote collective of transformation-based ensembles. *ACM TKDD*, 12:52, 2018.
- [16] Chang Wei Tan, François Petitjean, and Geoffrey I. Webb. Fasteer: Fast ensembles of elastic distances for time series classification. *DMKD*, 34:231–272, 2020.
- [17] Patrick Schäfer and Mikael Höggqvist. Sfa: A symbolic fourier approximation and index for similarity search in high dimensional datasets. In *EDBT*, pages 516–527, 2012.
- [18] Jessica Lin, Eamonn Keogh, Li Wei, and Stefano Lonardi. Experiencing SAX: A novel symbolic representation of time series. *DMKD*, 15:107–144, 2007.
- [19] Patrick Schäfer. The boss is concerned with time series classification in the presence of noise. *DMKD*, 29:1505–1530, 2015.
- [20] Patrick Schäfer and Ulf Leser. Multivariate time series classification with weasel+ muse. *arXiv:1711.11343*, 2017.
- [21] Matthew Middlehurst, James Large, Gavin Cawley, and Anthony Bagnall. The temporal dictionary ensemble (tde) classifier for time series classification. In *ECML/PKDD*, pages 660–676, 2020.
- [22] Houtao Deng, George Runger, Eugene Tuv, and Martyanov Vladimir. A time series forest for classification and feature extraction. *Information Sciences*, 239:142–153, 2013.
- [23] Matthew Middlehurst, James Large, and Anthony Bagnall. The canonical interval forest (CIF) classifier for time series classification. In *BigData*, pages 188–195. IEEE, 2020.
- [24] Nestor Cabello, Elham Naghizade, Jianzhong Qi, and Lars Kulik. Fast and accurate time series classification through supervised interval search. In *ICDM*, pages 948–953. IEEE, 2020.
- [25] Michael Franklin Mbouopda and Engelbert Mephu Nguifo. Scalable and Accurate Subsequence Transform. *hal-03087686v2*, 2021.
- [26] Ahmed Shifaz, Charlotte Pelletier, François Petitjean, and Geoffrey I Webb. Ts-chief: a scalable and accurate forest algorithm for time series classification. *DMKD*, 34:742–775, 2020.
- [27] Matthew Middlehurst, James Large, Michael Flynn, Jason Lines, Aaron Bostrom, Anthony Bagnall, Gustavo Batista, and Anthony Bagnall. Hive-cote 2.0: a new meta ensemble for time series classification. *Machine Learning*, 110:3211–3243, 2021.

- [28] Angus Dempster, François Petitjean, and Geoffrey I. Webb. Rocket: Exceptionally fast and accurate time series classification using random convolutional kernels. *DMKD*, 5:1454–1495, 2020.
- [29] Angus Dempster, Daniel F Schmidt, and Geoffrey I Webb. Minirocket: A very fast (almost) deterministic transform for time series classification. In *SIGKDD*, page 10. ACM, 2021.
- [30] M Gruber, T Dorst, A Schütze, S Eichstädt, and C Elster. *Advanced Mathematical and Computational Tools in Metrology and Testing XII*, chapter Discrete wavelet transform on uncertain data: Efficient online implementation for practical applications. 2020.
- [31] Haibo Liu, Ming Chen, Chong Du, Jiachang Tang, Chunming Fu, and Guilin She. A copula-based uncertainty propagation method for structures with correlated parametric p-boxes. *IJAR*, 2021.
- [32] F. Pedregosa et al. Scikit-learn: Machine learning in Python. *JMLR*, 12:2825–2830, 2011.
- [33] Emille E. O. Ishida. Machine learning and the future of supernova cosmology. *Nature Astronomy*, 3:680–682, 2019.
- [34] Kevin Fauvel, Elisa Fromont, Véronique Masson, Philippe Faverdin, and Alexandre Termier. Xem: An explainable-by-design ensemble method for multivariate time series classification. *DMKD*, 36(3):917–957, 2022.
- [35] Alejandro Pasos Ruiz, Michael Flynn, James Large, Matthew Middlehurst, and Anthony Bagnall. The great multivariate time series classification bake off: a review and experimental evaluation of recent algorithmic advances. *DMKD*, 35(2):401–449, 2021.
- [36] Marco Tulio Ribeiro, Sameer Singh, and Carlos Guestrin. " why should i trust you?" explaining the predictions of any classifier. In *KDD*, pages 1135–1144, 2016.
- [37] Scott M Lundberg and Su-In Lee. A unified approach to interpreting model predictions. In *NeurIPS*, pages 4765–4774. 2017.