

Fast Anytime Retrieval with Confidence in Large-Scale Temporal Case Bases

Mehmet Oğuz Mülâyim*, Josep Lluís Arcos

Artificial Intelligence Research Institute, IIIA-CSIC, Spanish National Research Council, Campus UAB, 08193, Barcelona, Spain

Abstract

This work is about speeding up retrieval in Case-Based Reasoning (CBR) for large-scale case bases (CBs) comprised of temporally related cases in metric spaces. A typical example is a CB of electronic health records where consecutive sessions of a patient forms a sequence of related cases. k-Nearest Neighbors (kNN) search is a widely used algorithm in CBR retrieval. However, brute-force kNN is impossible for large CBs. As a contribution to efforts for speeding up kNN search, we introduce an *anytime* kNN search methodology and algorithm. *Anytime Lazy kNN* finds exact kNNs when allowed to run to completion with remarkable gain in execution time by avoiding unnecessary neighbor assessments. For applications where the gain in exact kNN search may not suffice, it can be interrupted earlier and it returns best-so-far kNNs together with a confidence value attached to each neighbor. We describe the algorithm and methodology to construct a probabilistic model that we use both to estimate confidence upon interruption and to automatize the interruption at desired confidence thresholds. We present the results of experiments conducted with publicly available datasets. The results show superior gains compared to brute-force search. We reach to an average gain of 87.18% with 0.98 confidence and to 96.84% with 0.70 confidence.

Keywords: Large-Scale Case-Based Reasoning, Exact and Approximate k-Nearest Neighbor Search, Anytime Algorithms

1. Introduction

Industrial scale machine learning (ML) systems have to deal with larger amounts of digital data everyday due to the exponential growth of both its generation and availability [34]. Being a member of the instance-based learning subdivision of the larger ML family, many Case-Based Reasoning (CBR) systems are not exempt from this laborious opportunity either. Reminiscent of human thinking, CBR is based on two assumptions observed in real world that similar problems have similar solutions and that problems are likely to recur [27]. Hence, it stores past problem solving experiences as cases in its case base (CB) and when a new query is made to the system, it retrieves similar past problems in its CB and reuses their solutions by adapting them to the query [1, 24]. This type of reasoning is known as lazy learning in ML literature since a CBR system does not build a model prior to a query—as opposed to eager learning methods which do so—, and generalizes its cases every time a query is posed to the system. This behaviour is an advantage of CBR for continuously changing large CBs since it discards the need to re-train learned models with the updated data. However, due to its lazy nature, the efficiency of its retrieval phase affects overall system performance. And in practice, a growing CB eventually causes the so-called swamping utility problem [13, 38] which emerges when adding new cases to a CB degrades the system efficiency instead of improving it.

Being simple and effective, k-Nearest Neighbor (kNN) search is a widely used algorithm in CBR retrieval in particular and in instance-based learning in general. The naïve approach to find the kNNs of a query is to perform a brute-force search in the CB by evaluating the similarity of each case to the given query and return the k most similar cases. The runtime complexity¹ of this method may be acceptable for small sized CBs, but it implies an excessive

*Corresponding author

Email addresses: oguz@iia.csic.es (Mehmet Oğuz Mülâyim), arcos@iia.csic.es (Josep Lluís Arcos)

¹Runtime complexity of brute-force kNN is $O(ndk) + O(nk \log k)$ where n is the number of instances the query is made against, d the dimension of each instance, k the number of nearest neighbors searched for; the first part of the complexity is for distance calculations and the second part is for sorting the neighbors.

20 execution time for large-scale CBs due to expensive similarity calculations and is likely to evolve into above mentioned utility problem. There has been significant research to speed up nearest neighbor search (NNS) some of which we will review shortly in the next section. For occasions where the speed-up to find exact neighbors is still not computationally feasible, some efforts resorted to approximation methods to find approximate enough neighbors instead.

As a contribution to these efforts, to address both exact and approximate NNS, this article introduces an *anytime* kNN search algorithm, *Anytime Lazy kNN (ALK)*. We base our algorithm on a fast exact kNN algorithm *Lazy kNN* [33], and we extend it to a fully-fledged anytime algorithm. *ALK* finds exact kNNs when allowed to run to completion and otherwise, if interrupted, it returns best-so-far kNNs together with a *confidence* value attached to each neighbor. Confidence values reflect the expected qualities of approximate kNNs in terms of their similarities to the query compared to those of the exact kNNs. The proposed algorithm is also resumable and confidence values for approximate results increase over allocated time. Furthermore, we provide a means of confidence prediction to automatize the interruption of the algorithm by trading time with confidence in output.

30 *ALK* is equally efficient in exact kNN search as the original *Lazy kNN* in terms of avoiding unnecessary distance calculations that would be carried out by a standard brute-force search and it can save up significant execution time. Additionally, and as the main contribution of this work, we show that it reaches superior gains even when it is interrupted at very high confidence thresholds implying that we are very close to the exact kNNs. So, *ALK* gives the expert both the option to wait for the completion of the algorithm to obtain exact kNNs and the option to interrupt the search—manually and/or automatically—any time when a prompter response is needed and get best-so-far kNNs instead. In the latter case, he or she may also opt to resume the algorithm to get an output with a better confidence for the approximate results if so desires.

40 *Anytime Lazy kNN*, like its predecessor, excels specifically at domains where the CB can be organized as sequences of temporally related cases and the similarity metric takes into account the evolution of a sequence instead of treating each case individually. A good example to such domain can be found in *healthcare* where the electronic health record of a patient represents the sequence of his/her consecutive sessions and each new session is typically an update to this sequence. A search for similar patients regarding their medical histories should consider their session sequences. And, depending on whether the whole medical history or a part of it is queried, the similarity metric would use a time window encompassing the complete sequence or a subsequence of the health record respectively. Another natural example can be a *time series* (TS) dataset, where each instance is a sequence of temporally observed data. Here, each data point is essentially an update to the sequence and to assess the similarity between a query and a TS sequence in the dataset, the time window can cover the sequence fully or partially.

45 This article is organized as follows. Section 2 gives a background for our proposal. In subsection 2.1, first we briefly review the main approach in CBR community to overcome the utility problem, then we mention various methods developed so far to speed up NNS in instance-based learning in general. Subsection 2.2 clarifies the concepts we use throughout this paper describing the organization of a CB for our domain of interest. In subsection 2.3 we present information on the components and desired characteristics of anytime algorithms. We present the details of our proposed anytime kNN search methodology and our *Anytime Lazy kNN* algorithm itself in Section 3. Section 4 describes how we evaluated our algorithm and Section 5 gives highly encouraging results of experiments we conducted with real-world small to large time series datasets. Finally, we discuss the outcomes and future work in Section 6.

2. Background

2.1. Related work

60 Beside the obligation to deal with large scale data that comes with ever-growing CBs, current availability of the tools to interpret big data is also encouraging CBR researchers to work on systems that could benefit from hundreds of millions of cases (e.g. [19, 44]). Working with CBs of this scale could not be imagined until recently. Quite the contrary, till today the main approach in CBR community to tackle this problem has been to control the CB growth via case base maintenance (CBM) techniques while preserving the competence of the overall CB (for competence definition see [39], for CBM examples see [28, 20], for CBM dimensions see [43]). Of course, big data tools do not eliminate the need for CBM altogether, maintaining the correctness of cases and an adequate index are always useful for an efficient CBR. However, these tools may indeed outdate CB compression strategies as discussed in [19].

One of the efforts to overcome the computational overhead of kNN search has been to use parallel architectures (e.g. [3, 14, 23]). Other notable effort is using search trees such as k-d trees and its variants to partition the multidimensional search space and conduct NNS by pruning parts of the tree that cannot include the nearest neighbor (e.g. [42, 46], see also [22] for a comparison of search tree based NNS algorithms). However, beside the cost of their construction and maintenance, k-d trees are prone to the curse of dimensionality [5] phenomenon and although improvement suggestions exist (e.g. [12]), they are usually not recommended for high dimensional spaces².

Another common proposal to tackle the run-time bottleneck of NNS is using approximation in applications where approximate neighbors may be acceptable. Approximate methods retrieve a similar enough neighbor instead of the exact one in return for the improvement in retrieval speed. Typical techniques for approximate matching are using search tree adaptations (e.g. [31]), hashing (e.g. [2]) or proximity graphs (e.g. [17]). In CBR context, an approximate retrieval would imply approximate solutions. An example to an adaptation method of approximate solutions that could compensate for the retrieval accuracy in large case bases is suggested in [18].

When the retrieval time is a dire constraint, thanks to their ability to provide always a solution upon interruption, anytime algorithms [11, 48] have also been incorporated in NNS, such as in data stream mining (e.g. [26, 40]), databases (e.g. [45]) and in CBR retrieval [36] among other applications. If interrupted before completion, an anytime CBR retrieval would return the k best cases found so far with respect to their similarities to the query.

Ueno et al [40] proposes an anytime algorithm for nearest neighbor classification where they use a presorted (worst first) index which is created by assigning ranks to all instances based on their contribution in classification on the training set. This index is used as heuristics in NNS.

In Schaaf’s “Fish and Sink” (FaS) [36], the CB also holds “aspect” distances of the cases among each other which are weighted according to the “view” of the user asking the query. The NNS in CB starts with a predefined order of cases and directly tested (DT) cases “sink” with regards to their distances to the query, dragging down their view neighbors with them and labelling those neighbors as indirectly tested (IDT) cases. The time of interruption is important and only after all cases are labelled either as DT or IDT, FaS can show best k cases found so far regarding the relative depths of the DT and IDT cases. A prior interruption yields unconfident results. If FaS is not interrupted, it tests and sorts all cases.

Ricci and Avesani [35] use an anytime algorithm in learning a local similarity metric to be used in the CBR retrieval where the distance around a trial case is measured using the metric attached to that case. Their anytime algorithm updates the distance between an input case c and its neighbors depending on the role of the neighbors in solving c by incorporating a reinforcement learning procedure that adjusts the local weights.

On the other hand, with regard to the temporal dimension in CBR, Montani and Portinale [29] emphasize the importance of temporality especially in healthcare domain and propose a framework for the representation and retrieval of cases that are in the form of time series data in medical applications of CBR. The retrieval is addressed both at case and history levels, the latter taking into account the evolution of the temporally related cases.

As a more recent work, *Lazy kNN* [33] on which our algorithm is based on is a novel NNS which significantly reduces the number of evaluated cases in kNN search for large-scale CBs composed of temporally related cases. *Lazy kNN* leverages the triangle inequality in metric spaces by using it as a cutoff in NNS and evaluates only the “true kNN candidate” cases for a given query. A further advantage of the algorithm is that it does not need an extra data structure for an overall partitioning of the problem space, instead, it keeps track of the evaluated cases per sequence throughout the sequence’s history of updates.

2.2. Domain of interest

In CBR, a case typically consists of three parts: the *problem* representing the query, the *solution* suggested by the CBR system for the query and the *outcome* (feedback) after applying the solution. In this work, we only take into account problem parts of cases for CBR retrieval and use them as queries in kNN search. And we are particularly interested in domains where the similarity measure takes into account the history of updates to a sequence of temporally related cases instead of treating each update as an individual case.

For clarification purposes, the concepts *sequence*, *update*, *time window* and *query* that are used in this paper regarding our domain of interest are depicted in Figure 1.

²A good rule of thumb is to use k-d trees only if $N \gg 2^d$, where N is the number of data points and d is the number of dimensions [6].

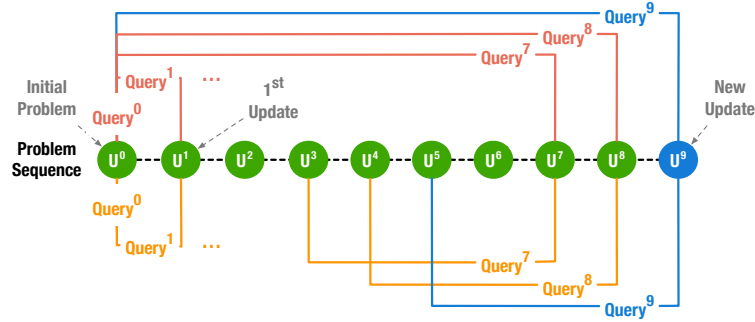


Figure 1: **Expanding (top) vs Fixed-width (below) Time Window approaches to form queries.** A problem *sequence* is created for each particular initial problem. Each consecutive problem related to this sequence is an *update*. A *query* is generated for each update, including the initial problem (i.e. 0^{th} update). To form the query for the u^{th} problem update; the former approach *expands* the *time window* encompassing all existing $u + 1$ updates; while the latter approach, for a *fixed-width* w , slides the window on the sequence covering the last w updates only (above $w=5$).

115 A problem sequence is created for each particular initial problem (e.g. a sequence is created for each patient in their first session) to represent the history of related problems to follow. We regard each consecutive problem related to this sequence as an update to the sequence. Consecutive queries for a problem sequence to be used in kNN search for CBR retrieval are formed as follows: The initial problem (which can be seen as the 0^{th} update) for every sequence is the first query. Each following update invokes the generation of a new query by applying a *time window* onto the sequence. In this work, we use *expanding* and *fixed-width* time window approaches. For each consecutive update, the former approach expands the time window encompassing whole problem sequence and the latter approach slides the time window of width w encompassing the last w updates to form the new query. Note that both approaches envelop the same updates for the first w problem updates, and for a given sequence, they generate the same number of queries (e.g. 10 queries in total in Figure 1).

125 In CBR context, after CBR's problem solving episode for a query, the query is saved as the problem part of a new case. For our domain, all consecutive episodes for the queries of the same problem sequence thus form a sequence of *temporally related cases*.

In implementation, case generation by time window approach can be conceptual or literal. In other words, there may be only a single data structure representing a problem sequence, and the case for a particular update can be formed on the fly when needed. Or, there may be l data structures for the cases of a sequence of length l .

130 The problem updates encompassed by the time window comprise the problem part of the case which is used by the similarity measure in kNN search of the CBR retrieval conducted for the query. Both approaches generate cases that encompass different number of updates. The CBR system designer should decide whether cases of different lengths should be compared in kNN search and if this is deemed necessary, the similarity metric—which has to satisfy the *triangle inequality* for our algorithm—should be implemented accordingly.

135 In the following subsection we highlight the important characteristics and concepts of anytime algorithms as a background for the proposed algorithm that excels in domains we have just described.

2.3. Anytime Algorithms

140 An anytime algorithm (AA) is a computer algorithm designed in such a way that given an input problem, it can provide a best-so-far solution at any time it is interrupted³ and the solution is accompanied by a quality value reflecting how close the interruption output is to the exact solution. In the core of a typical AA lies a function which incrementally improves the solution. The *quality measure*⁴ can be based on any characteristic of the output which is deemed important. It is preferable that the output quality monotonically increases over computation time and the improvement in quality is greater at the early stages of execution and it diminishes over time.

³Some anytime algorithms may need a short initialization time before they can be interrupted, e.g. [40].

⁴We use the terminology for anytime algorithm components given in [48].

145 After being interrupted, an AA can resume its execution if it is allocated more computation time. AAs also bear statistical information about their output quality over execution time for the input data they received. This information makes them predictable and it can be used for meta-reasoning about allocated computation time [16] (see [48] for a list of desired properties of AAs and quality metrics). All these characteristics make AAs useful in application domains where complete computation to solve the problem at hand is ineffective or impossible in real-time.

150 *Performance Profile* (PP) of an AA is used in estimating the quality of the output of the algorithm as a function of the amount of execution time [47]. An AA can have several PPs to track different result attributes [7]. After devising a quality measure, the PP of an anytime algorithm is typically constructed by using a simulation method. The algorithm is run with numerous input instances which are provided within a training dataset that can be randomly generated using the domain knowledge and the quality of the results over execution time are recorded. This statistical data composed of (*execution time, output quality*) pairs forms the *Quality Map* of the algorithm. Once the quality map is obtained, the corresponding PP of the algorithm can be derived from it as a formula by means of curve fitting methods or it can be represented as a table which reflects the discrete probability distribution of quality for discrete time allocations. The latter representation is called the *Performance Distribution Profile* (PDP) of the anytime algorithm and helps us give more accurate decisions compared to a single value of a fitted function. In our work, we opted for PDPs (see [15, 47] for further discussion on possible representations of PPs).

160 In accordance with the above description of a preferable AA, and regarding the focus of this work, an anytime kNN search algorithm is expected to monotonically improve on its k -nearest neighbors and provide a quality value attached to the best-so-far neighbors upon interruption. Also the improvement in output quality is preferred to be diminishing over time which means that the nearest neighbors found in early stages of search are almost as close to the query as the exact neighbors. Furthermore, the algorithm should be able to resume its execution without a major overhead if it is allocated extra computation time to improve on its results.

165 In the following section, after discussing the challenges in building such a well-tempered anytime kNN search algorithm, we point out why we deem Lazy kNN a very good candidate to be converted into an AA which overcomes these difficulties. Then, we provide the steps of this conversion as we detail the implementations of the above mentioned AA concepts at each step until we ultimately achieve our fully-fledged *Anytime Lazy kNN* search algorithm.

3. Anytime Lazy kNN

The main difficulty of converting an exact kNN search to an anytime algorithm lies in the quality assessment of the best-so-far neighbors. Having the search interrupted, we would like to compare the similarities of approximate and exact kNNs to the query. However, it is impossible to build an accurate quality measure for such an assessment. 175 The reason is obvious, exact kNNs remain unknown till the end of search and even though we might have already found them earlier, we cannot be aware of this until we evaluate all candidates in the search space. Consequently, for an algorithm which searches kNNs all at once (e.g. brute-force kNN search), exact kNNs are available only as a whole and after the completion of the search.

180 As an enhancement to partially ease the constraint of having to wait till completion, if we conduct kNN search in an incremental fashion finding k -nearest neighbors one by one in k iterations, we would at least guarantee the exactness of the top $i-1$ NNs when interrupted at the i^{th} iteration ($i \leq k$). Of course, this method would make sense only if extra iterations do not imply an additional cost of redundant similarity calculations. And clearly, although incremental search bears the possibility of providing some of the exact kNNs upon interruption, it does not eliminate the need to assess the quality of the remaining kNN list members that are yet-to-be ascertained for exactness.

185 In this section, we first deal with the design of an anytime kNN algorithm that exhibits the desired performance profile outlined in subsection 2.3. For the core of the AA, we propose to extend an existing algorithm, Lazy kNN [33]. Specifically, although Lazy kNN was designed to provide exact kNNs as a whole list after running to completion, we show how it can be converted to an incremental kNN search without any redundant computation in subsection 3.1. Then, we explain how Incremental Lazy kNN can be used in the core of an AA and present the pseudo-code of *Anytime Lazy kNN* (henceforth *ALK*) in subsection 3.2.

190 Once we have the AA algorithm, we introduce the steps required to define the appropriate mechanism to assess the quality of best-so-far kNNs at any given moment. In AA literature, when accuracy is not an option as a metric to build a deterministic quality measure, it is common to resort to a certainty metric to reflect a degree of correctness of intermediate results, e.g. by using the probability distribution of output quality over time [48].

195 Figure 2 illustrates the steps of the generation and use of the probabilistic quality measure that we implemented for our algorithm. Top two steps are to build the PDP of *ALK* for an application domain. In the first step, *ALK* runs kNN search simulations and gathers statistical data to generate the Quality Map of the algorithm. We detail this step in subsection 3.3. In the second step which is described in subsection 3.4, PDP is built out of the quality map. PDP endows our algorithm with the ability to predict the *expected quality* of best-so-far kNNs upon interruption. Henceforth, in concordance with CBR literature [9], we refer to expected output quality as *confidence* and define our confidence estimator based on PDP in subsection 3.5. We also show how PDP helps us to automatize the interruption upon reaching a given confidence threshold. In the final step of Figure 2, *ALK* is now ready to serve as an anytime kNN search algorithm. It accepts a query and an optional interruption point, and returns the exact/best-so-far kNNs together with a confidence value for each neighbor.

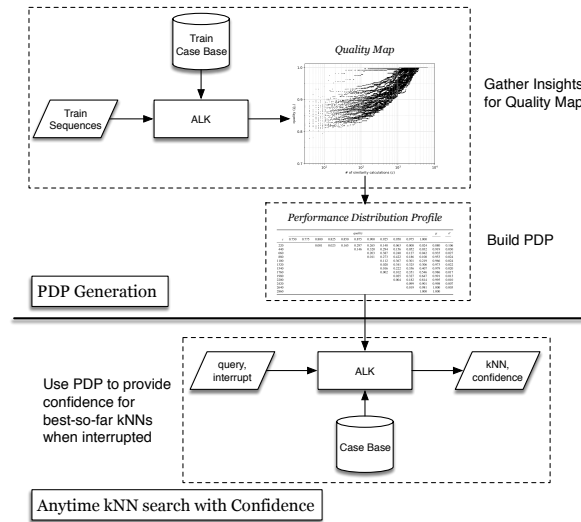


Figure 2: **Generation and use of confidence.** First, the *Quality Map* of the algorithm for the application domain is created out of interruption simulations with training data. Then, *Performance Distribution Profile* is generated out of *Quality Map*. PDP endows *ALK* with the ability to attach *confidence* values to its best-so-far kNNs when interrupted. PDP also provides a means to automatize interruption at given confidence thresholds.

205 3.1. Incremental Lazy kNN

Lazy kNN’s strength in speeding up exact kNN search comes from the evaluation of only the true kNN candidate cases in the CB for a query. During the execution of Lazy kNN, best-so-far kNN list members and their positions in the list are subject to change until the last candidate case in the CB is evaluated. And, this is only when the exact kNN list is provided as output. This is because, with respect to the triangular inequality used in the candidacy assessment, even the last candidate can potentially surpass the nearest neighbor found so far. Due to this nature of having to run to completion, Lazy kNN cannot say how confident it is of its best-so-far kNNs when interrupted.

To gradually improve on each of the kNNs and to be able to provide at least some of the exact NNs upon interruption, we developed an incremental version of Lazy kNN where we basically invoke the original algorithm *k* times iteratively. And, at each iteration *i* we find the *i*th exact NN. Thus, if the algorithm is interrupted during the *i*th iteration, it ensures that the top *i*-1 NNs are the exact NNs of the query.

The beauty of this conversion from standard to incremental kNN search is that, despite reiterations, Incremental Lazy kNN does not carry out any redundant similarity calculation compared to the original version of the algorithm. Though surprising it might be at first glance, this behaviour is due to the fact that we always evaluate the minimum number of candidate cases at each iteration, and after *k* iterations, the total number of assessed candidates equals the number of assessments made by the original Lazy kNN.

With respect to the desired monotonicity property of an AA, we can argue that exact kNN search in general exhibits monotonicity since any kNN search algorithm could maintain best-so-far neighbors even if it cannot improve any of them after a new neighbor candidate evaluation and provide these when interrupted. However, for all-at-once

Algorithm 1: Anytime Lazy KNN

```

1 Class AnytimeLazyKNN:
    Attributes: k: k of kNN
                _current_iter: Iteration index to start/resume the NNS from
                _current_index: Index to the candidate assessment in search space to start/resume the NNS from
                _query: The query formed for the problem update in the sequence for which the NNS is to be started/resumed
    Methods : _Construct (k):
                _IncrementalLazyKNN (query, interrupt)

2 Function _Construct (k):
3     this.k ← k
4     this._current_iter ← 1
5     this._current_index ← null
6     this._query ← null
7     return this

8 Function _IncrementalLazyKNN (this, query, interrupt):
9     if query <> this._query then // Update ?
10        this._query ← query
11        this._current_iter ← 1
12        this._current_index ← null
13    for iter ← this._current_iter to this.k do // Iterate ×k
14        kNN, this._current_index ← LazyKNN* (iter, this._query, interrupt, this._current_index)
15        if this._current_index <> null then // Interrupt ?
16            this._current_iter ← iter
17            break
18    return kNN, confidence (kNN, interrupt)

```

algorithms this gradual improvement is for the kNN list as a whole (i.e. any member can be improved during search any time). On the other hand, Incremental Lazy kNN possesses a monotonicity at individual nearest neighbor level that serves better for AA purposes. Given more time, current exact NNs will not change but the approximate ones are likely to be replaced by nearer neighbors, eventually all kNNs becoming the exact ones.

We note that the incremental nature of our algorithm is analogous to Broder’s incremental NNS [8] which also finds kNNs iteratively starting from the first, ending with the k^{th} ; but it differs from the incremental retrieval concept of Cunningham et al. [10] and Jurisica et al. [21] because there, the iteration takes place in conversational CBR systems where the retrieval is incrementally refined via iterative user interactions.

The following subsection gives our Anytime Lazy kNN search algorithm with Incremental Lazy kNN at its core.

3.2. Anytime Lazy kNN Algorithm

Besides monotonicity, another desired property of AAs is preemptability [48], that is, the capability to resume their execution after being interrupted. Incremental Lazy kNN can easily be made resumable by introducing two attributes to the algorithm to preserve the point in NNS where the interruption has occurred: 1) current iteration and, 2) the index of the next candidate to be assessed.

The simplified pseudo-code of *ALK* is given in Algorithm 1. For every particular problem sequence, a unique instance of the *AnytimeLazyKNN* class is instantiated (e.g. one instance for each patient). Then, each update (including the initial problem) to the sequence is passed as a *query* to the *IncrementalLazyKNN* method together with an optional interruption point (*interrupt*). *_current_iter* and *_current_index* are the two above-mentioned private instance attributes respectively to be used for resumability.

IncrementalLazyKNN uses a slightly extended version (*LazyKNN**) of the original Lazy kNN which accepts two additional optional arguments: *interrupt* and *_current_index* and returns new *_current_index* beside kNNs. *LazyKNN** is iteratively called k times, each time passing the iteration index *iter* as the k argument. The return value for *_current_index* is null if no interruption occurs, otherwise it points to the first candidate to be assessed if the algorithm is resumed.

Upon exit, *IncrementalLazyKNN* method returns the *kNN* list and the confidence of the algorithm for each member of the list. If the algorithm is run to completion, the kNNs will be exact and their confidence values will be 1. If the algorithm is interrupted, best-so-far kNN list is returned together with the expected quality values for each

member of the list provided by the `confidence` system function which is based on the PDP of the algorithm that we will explain in subsection 3.4.

But, as explained in 2.3, before building the PDP of *ALK*, we need to generate its quality map. And in the following subsection we show how we do it after defining our quality measure.

265 The complete code of *ALK* including all its functionality that will be covered throughout the rest of this article is publicly available at: <https://github.com/IIIA-ML/alk>

3.3. Quality Measure and Map

270 Quality measure of an AA is usually a function of execution time. However, in order to have a measure independent of the computer platform that the algorithm runs on, we opted to implement a measure which is a function of the number of similarity calculations carried out so far in kNN search. So, given a *query* and a number of calculations *c* as the interruption point, we define the quality measure for *ALK* as follows:

$$Q_c = \frac{\text{sim}(NN_c^k, \text{query})}{\text{sim}(NN_E^k, \text{query})} \quad (1)$$

where *sim* is the similarity metric, NN_c^k and NN_E^k are the best-so-far and exact k^{th} NN respectively, Q_c gives the output quality. If need be, *c* can be translated into actual execution time for the platform used.

265 Having the quality measure, the quality map can be generated as follows. After each similarity calculation *c* made for an input *query*, the similarity of the best-so-far k^{th} neighbor (i.e. the dividend in Eq. (1)) is recorded. When the simulation ends for that query, the similarity of the exact k^{th} NN (i.e. the divisor in Eq. (1)) is obtained. Then, by backtracking the simulation, the quality Q_c that we would get after each similarity calculation *c* is calculated using the recorded *sim* data. Eventually, all (c, Q_c) pairs collected during simulations provide us with the *QualityMap* of our algorithm.

270 However, this quality map does not reflect neither the temporal relations between cases nor the incremental nature of *ALK*. If we want a finer-grained quality measure incorporating these characteristics as well, we may add two more dimensions to the map. Since *ALK* finds the kNNs in an incremental fashion, the first extra dimension would be the index *i* of a nearest neighbor in the kNN list. This would allow us to assign a quality value per neighbor.

275 As for the temporal dimension, we could use the index *u* of the update on the problem sequence for which the query is generated for. This dimension provides even a finer-grained quality map, because, the more updates are covered by the time window, the less difference will have been introduced by the new update. Therefore, the more similar will be two successive queries and intuitively, the more similar will be their neighbors. Consequently, for a new query covering multiple updates, although *ALK* needs to evaluate all candidates within the neighbors of prior queries for the sake of mathematical exactness, it is very likely that the exact kNNs are found within the neighbors of the recent queries of the same sequence. As a result, this likelihood leads to higher quality values after fewer calculations. 280 This phenomenon also helps *ALK* to fulfill the desired AA property of diminishing output quality values. In other words, the increase in the quality of best-so-far kNNs are likely to be higher in the early similarity assessments during the kNN search compared to the later assessments.

Hence, we define the finer-grained quality as follows:

$$Q_c^{u,i} = \frac{\text{sim}(NN_c^{\mu,i}, \text{query}^u)}{\text{sim}(NN_E^{\mu,i}, \text{query}^u)} \quad (2)$$

285 where $NN_c^{\mu,i}$ is the i^{th} NN returned by the algorithm when it is interrupted after *c* similarity calculations during the NNS for the u^{th} query of a problem sequence, and $NN_E^{\mu,i}$ is the exact i^{th} NN for the same query.

290 Figure 3 shows an example to the quality map of *ALK* generated by using the quality measure in Eq. (2). It is generated throughout simulations with input sequences taken from the *SwedishLeaf* dataset (see subsection 4.2). The figure is a 2D excerpt of the 4D map plotted for the third nearest neighbor of the queries generated for the tenth updates (i.e. $Q_c^{10,3}$) of input sequences.

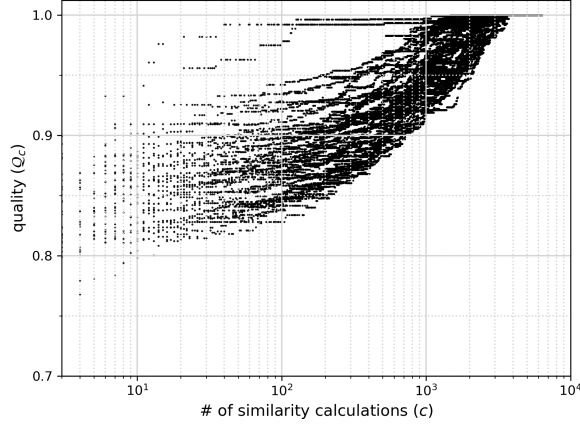


Figure 3: **Quality map** of Anytime Lazy kNN for the *SwedishLeaf* dataset. Here the map is shown for the 10th update and the 3rd nearest neighbor. Every point on the map is an instance of output quality observed throughout simulations.

3.4. Performance Distribution Profile

As the final step to have *Anytime Lazy kNN*, we use the quality map to build the PDP of our algorithm which is essentially the discrete probability distribution of output quality over computation time.

For this purpose, we first discretize the calculation range into m discrete calculation values c_1, \dots, c_m of equal intervals of δ_c , where c_m is the maximum number of similarity calculations performed for a test input during simulations. Then we discretize the quality range $[0, 1]$ into n discrete quality values q_1, \dots, q_n of equal intervals of δ_q .

In coherence with the finer-grained *QualityMap* presented above, we create a four-dimensional array \mathcal{PDP} to hold the discrete probability distribution of quality. The dimensions of \mathcal{PDP} are the maximum update number for an input sequence during simulations, k of kNN, the number of calculation intervals m and the number of quality intervals n respectively.

Thus, an entry in $\mathcal{PDP}[u, i, r, v]$ represents the discrete probability that the output quality of the best-so-far i^{th} NN is in $(q_v - \delta_q, q_v]$ after a number of similarity calculations in $(c_r - \delta_c, c_r]$ made during the kNN search for the u^{th} problem update.

The values of system parameters δ_c and δ_q , hence the size of the \mathcal{PDP} can be adjusted with respect to the desired accuracy of the performance information. The smaller δ_c and δ_q are, the more accurate will be the expected quality predicted by the \mathcal{PDP} . The corresponding *Performance Distribution Profile* of the quality map in Figure 3 is given in Table 1.

3.5. Confidence

Equipped with PDP, our algorithm becomes ready to predict the output quality when the kNN search for an unseen query is interrupted. We refer to the expected quality as the *confidence* of our algorithm in conformity with the CBR literature [9]. Specifically, we define our confidence measure as the weighted mean of the probability distribution of quality for a calculation interval in \mathcal{PDP} :

$$\mu_c^{u,i} = \sum_v q_v \mathcal{P}_c^{u,i}(q_v) \quad (3)$$

where $\mathcal{P}_c^{u,i}$ is a shorthand for $\mathcal{PDP}[u, i, r_c, v]$ and r_c is the interval in the calculation range where c falls into. This equation gives us the confidence $\in [0, 1]$ of the best-so-far i^{th} NN of the u^{th} update when *ALK* is interrupted after c number of similarity assessments during kNN search. Since confidence gives us a mean value μ , we also provide the (weighted) standard deviation σ of the quality distribution to be used together with it:

Table 1: **Performance Distribution Profile with Confidence.** The \mathcal{PDP} is essentially a 4-dimensional array; here the table for update=10 and the 3rd nearest neighbor in kNN entry corresponding to the quality map in Fig. 3 is shown. *confidence* for a calculation range is the expected quality which is defined as the weighted mean of the probability distribution of quality for that range. Note that $\delta_c=220$ and $\delta_q=0.025$ settings are used in the generation of this \mathcal{PDP} .

c	quality											μ	σ
	0.750	0.775	0.800	0.825	0.850	0.875	0.900	0.925	0.950	0.975	1.000		
220			0.001	0.023	0.165	0.297	0.265	0.140	0.063	0.008	0.024	0.880	0.106
440						0.146	0.320	0.294	0.156	0.052	0.032	0.919	0.030
660							0.203	0.387	0.240	0.127	0.042	0.935	0.027
880							0.011	0.273	0.422	0.186	0.108	0.953	0.024
1100								0.112	0.367	0.301	0.219	0.966	0.024
1320								0.028	0.341	0.325	0.306	0.973	0.022
1540								0.016	0.222	0.356	0.407	0.979	0.020
1760								0.002	0.102	0.351	0.546	0.986	0.017
1980									0.025	0.327	0.647	0.991	0.013
2200									0.004	0.182	0.814	0.995	0.010
2420										0.099	0.901	0.998	0.007
2640										0.019	0.981	1.000	0.003
2860											1.000	1.000	

$$\sigma_c^{u,i} = \sqrt{\sum_v (q_v - \mu_c^{u,i})^2 \mathcal{P}_c^{u,i}(q_v)} \quad (4)$$

The confidence μ and its deviation σ for the quality map example in Figure 3 are given in the Table 1. In this \mathcal{PDP} excerpt, we can see that σ is higher for early calculations when fewer candidates are assessed and it decreases as we assess more candidates. When c does not coincide exactly with the \mathcal{PDP} calculation intervals, we use linear interpolation for both μ and σ .

Beside giving us a confidence value to reason with the best-so-far kNNs, the PDP of our algorithm also provides us with a means to automatize the interruption itself. This may be achieved in two ways: either by specifying a time of execution or reaching a desired confidence threshold for the output.

In the former case, the execution time can be translated into a number of similarity calculations depending on the computer system the algorithm is running on. In the latter case, the PDP gives us the estimated number of similarity calculations to reach a desired confidence. In either case, obtained number of calculations is used for interruption.

While selecting interruption points, the standard deviation of the probability distribution of quality can be taken into account. Thus, if we define the *confidence threshold* for interruption as $\mu + z\sigma$, the choosing of the z offset parameter, that is the number of σ 's subtracted from or added to μ , would depend on how precautious or optimistic we want to be with the confidence provided by the \mathcal{PDP} respectively.

The trustworthiness of a probabilistic model depends on how much of the plausible problem space is represented by the model. In other words, the more representative the input queries used in simulations are for our domain, the more we can trust our model in predicting the confidence of our algorithm for future queries. Regarding CBR's fundamental representativeness assumption, we can safely assume that the CB that we will train our model with is representative of our problem space. Nevertheless, as we will describe in subsection 4.2, in our experiments we chose our training and test datasets in a way to enable a rigorous testing of the representativeness of the PDP as well.

4. Evaluation Methodology

There have been three main goals for the development of *ALK*: 1) To be able to interrupt kNN search and get best-so-far kNNs when exact kNN search is not feasible; 2) to attach confidence values to best-so-far kNNs that indicate how much we can trust each one of them in the reasoning process; and, 3) to be able to automatize interruption upon reaching given confidence thresholds. The previous section detailed the steps of how we developed such an anytime algorithm and a *confidence* measure that gives us the estimated output *quality*.

Confidence plays a key role in *ALK* both to assess the quality of the best-so-far output and to determine thresholds to automatize interruption. Therefore, to evaluate whether we met our above-mentioned goals, first we need to have a notion of efficiency for our confidence estimation. In other words, we would like to know how much we can trust




Figure 4: **Efficiency of confidence.** $\frac{\mu+z\sigma}{Q}$ ratio for interruption tests on the *SwedishLeaf* dataset, with $z=-1$ and different time window width w and $step$ settings and using the PDP for which an excerpt is given in Table 1.

the confidence measure itself. Accordingly, we define our efficiency measure in subsection 4.1 and explain how to interpret it. Later, in subsection 4.2, we describe the real world time series datasets that we used for our experiments to build CBs each of which was organized as sequences of temporally related cases. In subsection 4.3, we give the settings and insights of the experiments that we conducted both to demonstrate empirically that our confidence measure is efficient enough, and to provide evidence regarding the speed-up achieved by *ALK* when we interrupt the algorithm upon reaching given confidence thresholds.

4.1. Efficiency of confidence

Confidence μ is an estimator of how close the best-so-far kNNs are to the query compared with the exact kNNs. In other words, it is an estimator of the expected quality of the best-so-far kNNs. Therefore, to measure the efficiency of a confidence estimation, we utilize the ratio of the confidence μ and the observed actual quality Q and we formally define *efficiency* as follows:

$$\eta_c^{u,i} = \frac{\mu_c^{u,i} + z\sigma_c^{u,i}}{Q_c^{u,i}} \quad (5)$$

We use a $z\sigma$ offset value to be able to incorporate the deviation of the confidence provided by the *PDP* along with the μ itself upon interruption. The factor z can be chosen regarding how prudent we want to be with raw confidence values. $z=-1$ would be a more cautious choice than the neutral $z=0$, and $z=1$ would be a more optimistic one. An efficiency value $\eta \gg 1$ would signal an overconfident confidence measure that can possibly mislead reasoning with best-so-far kNNs; whereas $\eta \ll 1$ would imply an overcautious measure suggesting longer times of execution till reaching an acceptable approximation. On the other hand, while interpreting efficiency, we should bear in mind that due to its discrete nature, the precision of *PDP* (i.e. the choosing of δ_c and δ_q) affects the accuracy of quality estimation as well.

An example plot for the efficiency of confidence is shown in Figure 4. Interruption tests for this plot were conducted on four different CBs generated out of the *SwedishLeaf* dataset by four different w and $step$ settings. The confidence thresholds for interruption were selected with $z=-1$ setting. For all CBs, we see that the dispersion in efficiency is larger for lower confidence thresholds and it diminishes for higher thresholds. But, the efficiency almost consistently remains below 1.0, converging to 1.0 as the confidence threshold also gets closer to 1.0, to which we ultimately reach when we have the exact kNNs. This is a behavior that we desired by setting the z parameter to -1 preferring to be precautious with the confidence provided by the PDP. And indeed, Figure 4 shows that, when the

Table 2: **Time series datasets** used to generate CBs of temporally related cases in experiments for building PDPs and their corresponding CB sizes generated with two different time window *step* settings.

#	Dataset	Type	Sequences	Updates	CB Size	
					<i>step</i> =10	<i>step</i> =1
1	<i>PowerCons</i>	Device	180	144	2,700	25,920
2	<i>SwedishLeaf</i>	Image	625	128	8,125	80,000
3	<i>Strawberry</i>	Spectrograph	613	235	14,712	144,055
4	<i>EOGHorizontalSignal</i>	Medical	362	1,250	45,250	452,500
5	<i>InsectWingbeatSound</i>	Sensor	1,980	256	51,480	506,880
6	<i>ECG5000</i>	Medical	4,500	140	63,000	630,000
7	<i>UWaveGestureLibraryX</i>	Motion	3,582	315	114,624	1,128,330
8	<i>Yoga</i>	Image	3,000	426	129,000	1,278,000
9	<i>Phoneme</i>	Sound	1,896	1,024	195,288	1,941,504
10	<i>Mallat</i>	Simulated	2,345	1,024	241,535	2,401,280
11	<i>MixedShapesRegularTrain</i>	Image	2,425	1,024	249,775	2,483,200

algorithm was interrupted at $\mu\text{-}\sigma \in \{0.7, \dots, 0.98\}$, the actual quality of the best-so-far kNNs were higher than these confidence threshold values.

4.2. Datasets

375 We used eleven univariate time series datasets available at the UEA & UCR Time Series Classification Repository [4] as real world data in our experiments. Every dataset in the repository is available as a two-pack of train and test sub-datasets. Some test sub-datasets are larger than their train reciprocals. Since, for the scope of this article, we are only interested in the retrieval of the problem parts of cases and not in classification with their solutions, we took the liberty to use the larger one to generate the CB_TRAIN to be used in building the PDP for a dataset.

380 On the other hand, as in any probabilistic model, the representativeness of the PDP built for a specific training CB is crucial for the efficiency of the confidence estimation when the same PDP is used for future queries. Therefore, to be able to test the representativeness of the PDP as well, we carried out interruption tests on the CB_TEST generated out of the corresponding smaller sub-dataset. In other words, in interruption tests both the CB and the test sequences were unseen.

385 To build a CB out of a given univariate TS dataset, as described in subsection 2.2, we treated each instance in the dataset as a *sequence* for our CB and each data point of the TS instance as an *update* to the sequence. Then we applied *time window* on every instance to create individual *cases* of the sequence. Each window represented the problem part of a case and we regarded each data point enveloped by the window as a *feature* of that case. We adopted two approaches to implement time windows. The first approach used a sliding window of *fixed-width* and the second one an *expanding* window. For a TS instance of length l (i.e. having l data observations), both approaches generated l cases.

390 To generate even more diverse CBs out of a given TS dataset, we also incorporated time window *step* concept. With a *step* > 1 setting, we moved the time window in steps over the sequence and we generated cases for every *step* number of updates instead of doing it for each update in a sequence. Having *step* ≥ 1 as an optional parameter, for a problem sequence of length l , both expanding and fixed-window approaches generated $\lceil l/\text{step} \rceil$ number of cases. Table 2 summarizes the datasets and their corresponding CBs used in our experiments to build PDPs.

4.3. Experiment settings

400 With respect to distance measures, ALK is based on Lazy KNN algorithm that relies on distance measures that are true metrics as mentioned in subsection 2.1. Therefore, we used normalized *euclidean distance* in all experiments⁵. For every CB, normalization was done by taking into account the *min* and *max* values of the related univariate TS dataset. And the similarity *sim* between a *query* and a *case* in CB was defined as: $\text{sim}(\text{query}, \text{case}) = 1 - \text{distance}(\text{query}, \text{case})$.

⁵Interested reader can refer to [37, 41] for empirical comparison of similarity measures for time series data in general. Mueen et al.'s work [30] may be of special interest for time series where they also benefit from the triangle inequality for early abandoning of the costly distance measuring to find exact motifs.

As mentioned in subsection 2.2, regardless of the time window approach used, a decision has to be made with respect to how to measure the distance between two cases of different number of features. A straight-forward decision could be not measuring this distance at all and returning 0. However in order to test our algorithm with larger CBs, in our experiments, we opted to extend the shorter case to the length of the longer one by filling in missing features with values that maximized the distance.

After deciding the similarity assessment method, for each TS dataset given in subsection 4.2, we launched four experiments for combinations of two time window approaches (*Expanding* and fixed-width time window of width $w=40$) and two window step settings ($step \in [1, 10]$). Having k set to 9, each experiment for a configuration 3-tuple of (*dataset*, w , $step$) was conducted in three stages:

In the first stage, using the larger TS sub-dataset, we generated a set of sequences of temporally related cases for the experiment configuration. Then, we split this set into two parts; where one part served as the CB_TRAIN and the other part as input test sequences for our algorithm. For each input sequence, we generated input queries along its updates starting from the initial problem. By feeding the algorithm with these queries over CB_TRAIN, we generated the Quality Map for the dataset as described in subsection 3.3.

Consecutively, in the second stage, we built its PDP as described in subsection 3.4. While building PDP for each dataset, we discretized the range of similarity calculations and the quality range with $\delta_c = \lceil c_m / 400 \rceil$ and $\delta_q = 0.05$ interval settings respectively, where c_m was the highest number of similarity calculations reached for that experiment.

In the third stage, out of the smaller TS sub-dataset, we generated the CB_TEST and the set of unseen input queries for interruption tests, using the same method and time window configuration in the first stage. Then, we fed the algorithm with these queries over CB_TEST and for each query, we interrupted the algorithm using a set of confidence thresholds as interruption points.

In coherence with the dividend of the efficiency Eq. (5), confidence thresholds for interruptions were determined with a $z\sigma$ offset. We chose to be slightly cautious with PDP's confidence estimation and set z to -1 . The thresholds were chosen for the 9th NN and taking into account the problem update index of the query. So, for example, given the u^{th} query for a test sequence and a threshold of 0.95, the algorithm was automatically interrupted after the number of similarity calculations needed for the $\mu - \sigma$ provided by PDP for the 9th NN of a u^{th} query to reach 0.95.

5. Results

In this section, we provide the results of average gains in similarity assessments upon interruption and average efficiency of confidence estimation along experiments to assess if we met our design goals for *ALK*. Precisely, at each interruption, similarities of the best-so-far kNNs to the query, and the confidence μ together with its deviation σ for each member of the kNNs were recorded. Finally, we let the algorithm finish and we obtained the similarities of the exact kNNs to the query. And later, by backtracking, we calculated the actual qualities Q . This allowed us to calculate and record the efficiency η of the confidence for each NN using the Eq. (5). At interruptions, we also recorded the gain in similarity calculations for the query, i.e. the percentage of the avoided calculations upon interruption with respect to the total number of calculations that would have been carried out by a brute-force search.

Finally, recorded efficiency and gain values throughout experiments gave us the answers we were looking for. We show the average gain of the algorithm for the set of confidence thresholds used as interruption points in Table 3. In the same table, we also provide the average efficiency of confidence for each experiment together with its average deviation.

The average gain at an interruption threshold was calculated out of the gains of all test queries at that threshold. While interpreting the average gains, we note that, especially for the expanding window setting, the gain for later updates will be greater than earlier updates of a sequence. As mentioned in subsection 3.1, gains of uninterrupted Anytime Lazy kNN are precisely the gains that would be achieved by the original Lazy kNN. In Table 3, average gains for uninterrupted runs are in the range of [28.07, 96.35]. While the upper limit of this range can correspond to quite acceptable execution times for some CBs to wait for the exact kNNs, the lower limit may not be tolerable for very large CBs. And in the latter case, we would have to trade time for approximate results and this is when we benefit from the true merit of *ALK*.

In the table, for many configurations, we observe a notable leap between the gain for an uninterrupted run and the corresponding gain upon an interruption at a confidence threshold as high as 0.98. And for some experiment

Table 3: **Average Gain % upon Interruptions at Confidence Thresholds.** The algorithm is interrupted at the number of calculations when the $\mu\text{-}\sigma$ ($z=-1$) value provided by \mathcal{PDP} reaches to the given confidence thresholds. For e.g., a threshold of 0.95 means “stop when $\mu\text{-}\sigma$ reaches 0.95”. The table summarizes average *gains* in terms of % of avoided similarity assessments during tests compared to a brute-force search. The gains upon interruption can be compared to the gain at the ‘Unint.’ column which is achieved when the algorithm is run to completion uninterrupted for exact kNNs. We also provide the mean of efficiency η together with the mean of its standard deviation σ_η to reflect the *efficiency* of the quality estimation (i.e. *confidence*) for each experiment.

Dataset	Time window		Unint.	Interruption at $\mu\text{-}\sigma =$							Efficiency		
				0.98	0.95	0.92	0.90	0.85	0.80	0.75	0.70	$\bar{\eta}$	$\bar{\sigma}_\eta$
	<i>w</i>	<i>step</i>											
<i>PowerCons</i>	Expanding	1	75.65	97.78	98.63	98.96	99.01	99.17	99.30	99.38	99.45	0.92	0.07
		10	47.67	75.47	84.37	87.62	88.06	89.46	91.10	92.27	93.47	0.95	0.07
	40	1	69.62	97.44	98.49	98.91	98.96	99.16	99.27	99.39	99.46	0.92	0.08
		10	35.12	74.78	81.52	84.14	84.32	85.33	88.91	91.26	96.04	0.97	0.05
		1	75.96	98.79	99.14	99.25	99.31	99.43	99.52	99.61	99.66	0.92	0.07
		10	47.55	76.04	79.95	81.38	82.10	87.00	91.68	93.24	96.36	0.94	0.07
<i>SwedishLeaf</i>	Expanding	1	79.11	98.97	99.14	99.25	99.32	99.44	99.54	99.62	99.66	0.92	0.07
		10	38.01	63.69	75.34	84.17	89.44	95.32	96.29	96.34	96.37	0.94	0.07
	40	1	81.59	98.74	99.10	99.20	99.26	99.39	99.50	99.55	99.57	0.92	0.06
		10	56.39	74.92	80.33	83.01	83.51	87.07	96.16	96.75	97.07	0.94	0.06
		1	89.08	98.95	99.15	99.25	99.30	99.43	99.52	99.55	99.57	0.93	0.06
		10	57.79	79.98	83.81	85.27	85.91	92.04	95.29	98.03	98.34	0.94	0.07
<i>EOGHorizontalSignal</i>	Expanding	1	91.12	99.51	99.60	99.68	99.72	99.73	99.73	99.73	99.73	0.95	0.03
		10	76.32	98.84	98.96	99.02	99.20	99.41	99.50	99.57	99.63	0.94	0.07
	40	1	96.35	99.39	99.47	99.63	99.73	99.73	99.73	99.73	99.73	0.95	0.03
		10	87.04	99.08	99.21	99.24	99.24	99.25	99.29	99.33	99.63	0.91	0.08
		1	81.50	95.22	96.09	96.90	97.18	97.34	97.43	97.48	97.50	0.94	0.05
		10	55.96	58.04	62.29	68.29	71.22	78.89	87.11	91.28	93.59	0.89	0.09
<i>InsectWingbeatSound</i>	Expanding	1	84.97	95.40	96.25	97.03	97.25	97.41	97.50	97.56	97.60	0.95	0.05
		10	42.27	44.81	48.76	61.17	70.70	84.19	88.62	90.03	94.11	0.91	0.08
	40	1	76.71	92.23	94.56	95.89	96.42	97.09	97.32	97.41	97.46	0.93	0.05
		10	48.09	48.09	48.09	48.09	48.09	65.35	86.86	87.28	95.64	0.94	0.07
		1	79.90	93.51	95.34	96.37	96.61	97.07	97.31	97.40	97.45	0.94	0.05
		10	48.79	73.24	81.79	85.49	85.63	86.00	86.37	86.74	89.11	0.92	0.08
<i>ECG5000</i>	Expanding	1	84.20	98.15	98.66	98.84	98.87	98.91	98.93	98.95	98.95	0.96	0.03
		10	61.93	84.27	87.09	89.42	90.60	93.00	95.62	97.76	98.06	0.91	0.07
	40	1	91.78	98.11	98.75	98.89	98.91	98.95	98.97	98.99	98.99	0.96	0.03
		10	61.62	90.84	93.73	93.90	94.00	94.26	97.69	97.76	98.13	0.93	0.06
		1	86.70	95.22	96.65	97.04	97.14	97.22	97.24	97.26	97.28	0.96	0.03
		10	67.32	79.12	85.34	87.85	88.74	90.99	94.15	95.65	96.38	0.89	0.08
<i>Yoga</i>	Expanding	1	92.59	95.42	96.66	96.89	96.93	97.20	97.23	97.26	97.28	0.96	0.03
		10	63.99	70.68	89.03	91.44	91.83	92.84	95.20	95.70	95.79	0.93	0.06
	40	1	88.95	96.27	97.54	97.56	97.56	97.56	97.57	97.57	97.57	0.97	0.02
		10	71.60	90.53	93.00	94.48	95.25	96.42	97.06	97.37	97.52	0.92	0.06
		1	66.97	96.29	97.54	97.56	97.57	97.57	97.57	97.57	97.57	0.97	0.02
		10	28.07	94.06	94.63	95.13	95.46	96.25	97.27	97.31	97.50	0.94	0.05
<i>Phoneme</i>	Expanding	1	89.66	90.29	90.66	90.75	90.78	90.82	90.84	90.87	90.87	0.97	0.02
		10	72.83	73.29	74.24	76.14	77.86	83.29	86.10	86.84	86.96	0.91	0.06
	40	1	93.07	93.48	93.98	94.07	94.13	94.21	94.25	94.28	94.30	0.97	0.02
		10	67.74	69.40	71.43	75.86	77.85	83.05	85.49	86.00	86.97	0.93	0.05
		1	90.52	97.92	98.60	98.63	98.65	98.66	98.67	98.67	98.67	0.97	0.02
		10	74.30	96.23	96.94	97.32	97.55	98.06	98.38	98.51	98.59	0.92	0.06
<i>MixedShapesRegularTrain</i>	Expanding	1	95.68	97.95	98.56	98.57	98.62	98.67	98.68	98.68	98.68	0.98	0.02
		10	76.13	95.36	95.70	97.11	97.32	97.62	97.63	98.43	98.56	0.95	0.04

configurations, like the *Phoneme* with $w=40$ and $step=10$, this difference is tremendous. In this example, although the algorithm reaches a confidence threshold of 0.98 with an average 94.06% gain, it ends up doing many more calculations to ascertain the exact kNNs which ultimately reduces the gain down to 28.07% level. The confidence being quite efficient ($\bar{\eta}=0.94, \bar{\sigma}_\eta=0.05$) for this experiment, this phenomenon occurs due to the fact that, in this particular CB, there are many similar kNN candidates which need to be assessed, but, in the end they cannot win over best-so-far kNNs. Another common observation is that we reach higher gains for $step=1$ compared to $step=10$, because with the former configuration, less change is introduced per update and the kNNs of the new query are more similar to the kNNs of the previous one. Thus, less calculations are needed to obtain the exact kNNs and, in the case of interruption, to reach the desired confidence threshold.

We also see that for $z=-1$ setting, the quality estimation of PDP (i.e. the confidence) was quite efficient with a relatively minor deviation ($\bar{\eta} \in [0.89, 0.98], \bar{\sigma}_\eta \in [0.02, 0.09]$) throughout our experiments. The average efficiency was

Table 4: **Average Execution Times and Speed-up Factors per Update.** Average execution times (in milliseconds) of kNN search per sequence update are given for select experiments with time window width $w=40$ setting. Speed-up factors are given with respect to the execution time of a brute-force search. Note that, average execution time and speed-up factor for an experiment are proportionate to its average gain in Table 3.

Dataset	step	Uninterrupted			ALK interrupted at $\mu-\sigma =$							
		Brute		ALK	0.98		0.95		0.85		0.70	
		ms	ms	Speed	ms	Speed	ms	Speed	ms	Speed	ms	Speed
<i>PowerCons</i>	1	425.86	129.37	$\times 3.29$	10.90	$\times 39.06$	6.43	$\times 66.23$	3.58	$\times 119.05$	2.30	$\times 185.19$
	10	24.10	15.64	$\times 1.54$	6.08	$\times 3.97$	4.45	$\times 5.41$	3.54	$\times 6.82$	0.95	$\times 25.25$
<i>SwedishLeaf</i>	1	1295.09	270.54	$\times 4.79$	13.34	$\times 97.09$	11.14	$\times 116.28$	7.25	$\times 178.57$	4.40	$\times 294.12$
	10	69.93	43.28	$\times 1.62$	25.39	$\times 2.75$	17.25	$\times 4.06$	3.27	$\times 21.37$	2.54	$\times 27.55$
<i>Strawberry</i>	1	1437.21	156.94	$\times 9.16$	15.09	$\times 95.24$	12.22	$\times 117.65$	8.19	$\times 175.44$	6.18	$\times 232.56$
	10	98.52	41.59	$\times 2.37$	19.72	$\times 5.0$	15.95	$\times 6.18$	7.84	$\times 12.56$	1.64	$\times 60.24$
<i>EOGHorizontalSignal</i>	1	8057.42	294.10	$\times 27.4$	49.15	$\times 163.93$	42.70	$\times 188.68$	21.76	$\times 370.37$	21.76	$\times 370.37$
	10	710.63	92.10	$\times 7.72$	6.54	$\times 108.7$	5.61	$\times 126.58$	5.05	$\times 140.85$	2.27	$\times 312.5$
<i>InsectWingbeatSound</i>	1	958.27	144.03	$\times 6.65$	44.08	$\times 21.74$	35.94	$\times 26.67$	24.82	$\times 38.61$	23.00	$\times 41.67$
	10	54.77	31.62	$\times 1.73$	30.22	$\times 1.81$	28.06	$\times 1.95$	8.66	$\times 6.33$	3.23	$\times 16.98$
<i>ECG5000</i>	1	1370.28	275.43	$\times 4.98$	88.93	$\times 15.41$	63.86	$\times 21.46$	40.15	$\times 34.13$	34.94	$\times 39.22$
	10	69.98	35.84	$\times 1.95$	18.73	$\times 3.74$	12.74	$\times 5.49$	9.80	$\times 7.14$	7.62	$\times 9.18$
<i>UWaveGestureLibraryX</i>	1	6408.30	526.76	$\times 12.17$	121.12	$\times 52.91$	80.10	$\times 80.0$	66.01	$\times 97.09$	64.72	$\times 99.01$
	10	354.34	136.00	$\times 2.61$	32.46	$\times 10.92$	22.22	$\times 15.95$	8.19	$\times 43.29$	6.63	$\times 53.48$
<i>Yoga</i>	1	2329.18	172.59	$\times 13.5$	106.68	$\times 21.83$	77.79	$\times 29.94$	64.52	$\times 36.1$	63.35	$\times 36.76$
	10	138.37	49.83	$\times 2.78$	40.57	$\times 3.41$	15.18	$\times 9.12$	6.64	$\times 20.83$	5.83	$\times 23.75$
<i>Phoneme</i>	1	4383.57	1447.89	$\times 3.03$	162.63	$\times 26.95$	107.84	$\times 40.65$	106.52	$\times 41.15$	106.52	$\times 41.15$
	10	435.54	313.28	$\times 1.39$	25.87	$\times 16.84$	23.39	$\times 18.62$	11.89	$\times 36.63$	10.89	$\times 40.0$
<i>Mallat</i>	1	867.53	60.12	$\times 14.43$	56.56	$\times 15.34$	52.23	$\times 16.61$	49.88	$\times 17.39$	49.45	$\times 17.54$
	10	62.13	20.04	$\times 3.1$	19.01	$\times 3.27$	17.75	$\times 3.5$	9.02	$\times 6.89$	8.10	$\times 7.67$
<i>MixedShapesRegularTrain</i>	1	9590.46	414.31	$\times 23.15$	196.60	$\times 48.78$	138.10	$\times 69.44$	126.59	$\times 75.76$	126.59	$\times 75.76$
	10	1076.32	256.92	$\times 4.19$	49.94	$\times 21.55$	46.28	$\times 23.26$	25.51	$\times 42.19$	15.50	$\times 69.44$

Test environment specs: CPU: 12 \times Intel(R) Core(TM) i7-8700K @ 3.70GHz; Memory: 31GiB; OS: Ubuntu 18.04; Python: 3.7.3.

slightly below 1, which means that the actual quality was a bit above the confidence threshold. This was a desired behaviour of efficiency, since we wanted to be precautious by lowering the confidence μ by $-\sigma$ for interruptions. In other words, we preferred to have slightly lower gain in execution time to giving slightly overconfident results.

The *gain* concept that we have been using throughout the article is based on the percentage of the avoided similarity calculations compared to the number of similarity calculations that would have been carried out by brute-force search. This definition let us establish a platform-independent measure for the speed-up in kNN search by *ALK*. To give a more thorough view of this speed-up, we also provide Table 4 that translates the gain to real execution time for a test platform that we used. The table gives the average execution times for kNN search per sequence update conducted by brute-force search and *ALK*. The former evaluates all the cases in the case base while our algorithm assesses only the true kNN candidates. The table shows that even when *ALK* is not interrupted and run to completion to find exact kNNs, it is faster than brute-force search by orders of magnitude. However, the real contribution of our algorithm is for the occasions when this speed-up is still not feasible and we have to resort to approximate kNNs. In this case, even when we interrupt *ALK* at a high confidence threshold like 0.98, the speed-up drastically increases. For e.g., for the *EOGHorizontalSignal* experiment with $w=40$ and $step=1$ setting, *ALK* delivers best-so-far kNNs of the query for a sequence update with 0.98 confidence 163.93 times faster than the brute-force search on average, reducing the execution time from 8.057 seconds down to 0.049 seconds. For the same experiment, the speed-up factor increases to 370.37 if we settle with a confidence of 0.85. Table 4 also reveals that the speed-up is more dramatic for the larger case base of the same dataset. The CB generated for a dataset with $step=1$ is 10 times larger than the CB of a $step=10$ as explained in subsection 4.2. And this observation underpins the very purpose of the *ALK*, larger the CB is, higher becomes the speed-up. We also note that the average execution time and speed-up factor for an experiment are inversely proportionate to the corresponding average gain given in Table 3.

For each application, the definition of acceptable approximate results will be different. For some critical decisions, we may need approximate results with very high confidence whereas for less critical situations we may conform with less confidence. The gain that we will achieve for interrupting the algorithm with higher or lower thresholds will change according to the nature of the CB and used time window configurations. But in any case, if we opt for the approximate results instead of the exact ones, *ALK* will boost the speed-up in kNN search higher than Lazy kNN.

6. Discussion and Conclusions

490 In this work we introduced *Anytime Lazy kNN (ALK)*, an anytime exact and approximate kNN algorithm to boost CBR retrieval in large-scale CBs of temporally related cases, e.g. a CB of electronic health records of patients. Our algorithm is based on Lazy kNN [33] which is an effective exact kNN algorithm in CBR literature for such domains. However, for some applications, the notable speed-up provided by this algorithm may not suffice and the execution time for exact kNN may still be intolerable. Therefore, to fit the CBR system with an approximate retrieval
495 option for time-critical applications, we described our methodology to transform Lazy kNN to *ALK*. Specifically, after presenting our domain of interest and the concepts we used throughout this paper, we detailed how we constructed a probabilistic model of adjustable accuracy to estimate the quality of best-so-far kNNs upon interruption. We referred to the expected quality as the *confidence* of the algorithm for its output in accordance with the CBR literature. Later, we showed how we can implement confidence thresholds to automatize the interruption with options to be precautionous,
500 neutral or optimistic about the confidence provided by our probabilistic model. Before experimentation, we explained how we can treat a time series as a sequence of temporally related cases and showed different ways of generating CBs out of a TS dataset by different time window settings. Furthermore, we devised a means to measure the *efficiency* of confidence estimation to be used throughout experiments. Finally, we presented the results of numerous experiments conducted on publicly available TS datasets of diverse domains and characteristics.

505 The results show us that we can reach superior speed-up in approximate kNN search even when we interrupt the algorithm at very high confidence thresholds which means that best-so-far kNNs are almost as near to the query as the exact kNNs. So, with *ALK*, the expert can opt to wait for the completion of the algorithm to obtain exact kNNs or, he or she can interrupt the search manually/automatically any time when a prompter response is needed and get best-so-far kNNs together with a confidence value for each NN to reason with.

510 The *Performance Distribution Profile* used by our algorithm as the basis of its probabilistic model discretizes the calculation range linearly. A way to increase both the accuracy and efficiency of confidence can be to discretize this range logarithmically, attributing more importance to the beginning of the calculation range given the fact that NNs are actually found in the early stages of *ALK* execution for the domains of our interest.

515 We mainly focused on its use in a CBR context, but *ALK* can be applied to any kNN search domain that exhibits temporal relation between examples in a metric search space. For example, we tested *ALK* with univariate TS datasets, but in a like manner, given an adequate metric distance for a multivariate TS (see [25] for examples), our algorithm would work just the same. One way to build a CB from a multivariate TS dataset can be to treat each instance as a sequence and each time-dependent variable of the instance as a multi-valued case feature. The number of feature values would be determined by the applied time window.

520 On the other hand, we only used the problem space for retrieval. However, a case also bears a solution part and as future work we want to enhance retrieval by incorporating the solution space as well. Especially, low confidence zones in a CB's solution space (e.g. categorized as dubiousity patterns in [32]) can help *ALK* further tune its retrieval performance. For example, if the solutions of best-so-far kNNs exhibit a "border" pattern where two competing solutions (say classes) exist; then the algorithm may use this information to encourage resuming the search expecting
525 a possible change in kNNs favouring one of the competing solutions.

The complete code of *ALK* algorithm including all its functionality covered throughout this article and simple instructions to reproduce the experiments are publicly available at the online repository: <https://github.com/IIIA-ML/alk>

Acknowledgements

We thank to all contributors and maintainers of the UEA & UCR Time Series Classification Repository. This work has been funded by the project Playing and Singing for the Recovering Brain: Efficacy of Enriched Social-Motivational Musical Interventions in Stroke Rehabilitation (Play&Sing), 201729.31, Fundació La Marató de TV3; and, by the project Innobrain, COMRDI-151-0017 (RIS3CAT comunitats), and Feder funds. Mehmet Oğuz Mülâyim is a PhD Student of the doctoral program in Computer Science at the Universitat Autònoma de Barcelona.

References

- [1] Agnar, A., Plaza, E., 1994. Case-Based Reasoning: Foundational Issues, Methodological Variations, and System Approaches. *AI Communications* 7, 39–59. doi:10.3233/AIC-1994-7104.
- [2] Andoni, A., Indyk, P., 2006. Near-Optimal Hashing Algorithms for Approximate Nearest Neighbor in High Dimensions, in: *IEEE Symposium on Foundations of Computer Science*, IEEE. pp. 459–468. doi:10.1109/FOCS.2006.49.
- [3] Arefin, A.S., Riveros, C., Berretta, R., Moscato, P., 2012. GPU-FS-kNN: A Software Tool for Fast and Scalable kNN Computation Using GPUs. *PLoS ONE* 7. doi:10.1371/journal.pone.0044000.
- [4] Bagnall, A., Lines, J., Vickers, W., Keogh, E., . The UEA & UCR Time Series Classification Repository. URL: <http://www.timeseriesclassification.com>. (Last accessed 20 January 2020).
- [5] Bellman, R., 1957. *Dynamic Programming*. 1 ed., Princeton University Press, Princeton, NJ, USA.
- [6] Bentley, J.L., 1975. Multidimensional binary search trees used for associative searching. *Communications of the ACM* 18, 509–517. doi:10.1145/361002.361007.
- [7] Boddy, M., Dean, T., 1989. Solving Time-Dependent Planning Problems, in: *Eleventh International Joint Conference on Artificial Intelligence*, pp. 979–984.
- [8] Broder, A.J., 1990. Strategies for efficient incremental nearest neighbor search. *Pattern Recognition* 23, 171–178. doi:10.1016/0031-3203(90)90057-R.
- [9] Cheetham, W., 2000. Case-Based Reasoning with Confidence, in: *European Workshop on Advances in Case-Based Reasoning*, Springer. pp. 15–25. doi:10.1007/3-540-44527-7_3.
- [10] Cunningham, P., Smyth, B., Bonzano, A., 1998. An incremental retrieval mechanism for case-based electronic fault diagnosis. *Knowledge-Based Systems* 11, 239–248. doi:10.1016/S0950-7051(97)00049-X.
- [11] Dean, T., Boddy, M., 1988. An Analysis of Time-Dependent Planning, in: *Seventh AAAI National Conference on Artificial Intelligence*, AAAI. pp. 49–54.
- [12] Eastman, C.M., Weiss, S.F., 1982. Tree structures for high dimensionality nearest neighbor searching. *Information Systems* 7, 115–122. doi:10.1016/0306-4379(82)90023-0.
- [13] Francis, A.G., Ram, A., 1993. The Utility Problem in Case-Based Reasoning, in: *AAAI Case-Based Reasoning Workshop*, AAAI Press, Washington DC. p. 160.
- [14] Garcia, V., Debreuve, E., Barlaud, M., 2008. Fast k nearest neighbor search using GPU. *IEEE Computer Society Conference on Computer Vision and Pattern Recognition Workshops, CVPR Workshops*, 1–6doi:10.1109/CVPRW.2008.4563100.
- [15] Grass, J., 1996. Reasoning about computational resource allocation. *XRDS: Crossroads, The ACM Magazine for Students* 3, 16–20. doi:10.1145/332148.332154.
- [16] Grass, J., Zilberstein, S., 1996. Anytime algorithm development tools. *ACM SIGART Bulletin* 7, 20–27. doi:10.1145/242587.242592.
- [17] Hajebi, K., Abbasi-Yadkori, Y., Shahbazi, H., Zhang, H., 2011. Fast Approximate Nearest-neighbor Search with K-nearest Neighbor Graph, in: *International Joint Conference on Artificial Intelligence*, AAAI Press. pp. 1312–1317. doi:10.5591/978-1-57735-516-8/IJCAI11-222.
- [18] Jalali, V., Leake, D.B., 2017. Scaling Up Ensemble of Adaptations for Classification by Approximate Nearest Neighbor Retrieval, in: *International Conference on Case-Based Reasoning*, Springer, Berlin, Heidelberg. pp. 154–169. doi:10.1007/978-3-319-61030-6_11.
- [19] Jalali, V., Leake, D.B., 2018. Harnessing Hundreds of Millions of Cases: Case-Based Prediction at Industrial Scale, in: *International Conference on Case-Based Reasoning*, Springer Verlag. pp. 153–169. doi:10.1007/978-3-030-01081-2_11.
- [20] Juarez, J.M., Craw, S., Lopez-Delgado, J.R., Campos, M., 2018. Maintenance of case bases: Current algorithms after fifty years, in: *IJCAI International Joint Conference on Artificial Intelligence*, pp. 5457–5463. doi:10.24963/ijcai.2018/770.
- [21] Jurisica, I., Glasgow, J., Mylopoulos, J., 2000. Incremental iterative retrieval and browsing for efficient conversational CBR systems. *Applied Intelligence* 12, 251–268. doi:10.1023/A:1008375309626.
- [22] Kibriya, A.M., Frank, E., 2007. An Empirical Comparison of Exact Nearest Neighbour Algorithms, in: *Knowledge Discovery in Databases: PKDD 2007*. Springer Berlin Heidelberg, pp. 140–151. doi:10.1007/978-3-540-74976-9_16.
- [23] Kolodner, J., 1988. Retrieving events from a case memory: A parallel implementation. *Proc. of 1988 Case-Based Reasoning*, 233–249.
- [24] Kolodner, J., 1993. *Case-Based Reasoning*. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA.
- [25] Kotsakos, D., Trajcevski, G., Gunopulos, D., Aggarwal, C.C., 2013. Time-Series Data Clustering, in: Aggarwal, C.C., Reddy, C.K. (Eds.), *Data Clustering*. Chapman and Hall/CRC. chapter 15, pp. 357–380. doi:10.1201/9781315373515-15.
- [26] Kranen, P., Seidl, T., 2009. Harnessing the strengths of anytime algorithms for constant data streams. *Data Mining and Knowledge Discovery* 19, 245–260. doi:10.1007/s10618-009-0139-0.
- [27] Leake, D.B., 1996. CBR in Context: The Present and Future, in: Leake, D.B. (Ed.), *Case-Based Reasoning: Experiences, Lessons & Future Directions*. The AAAI Press/The MIT Press, pp. 3–30.
- [28] Leake, D.B., Smyth, B., Wilson, D.C., Yang, Q., 2001. Introduction to the Special Issue on Maintaining Case-Based Reasoning Systems. *Computational Intelligence* 17, 193–195. doi:10.1111/0824-7935.00139.
- [29] Montani, S., Portinale, L., 2006. Accounting for the temporal dimension in case-based retrieval: A framework for medical applications. *Computational Intelligence* 22, 208–223. doi:10.1111/j.1467-8640.2006.00284.x.
- [30] Mueen, A., Keogh, E., Zhu, Q., Cash, S., Westover, B., 2009. Exact Discovery of Time Series Motifs, in: *SIAM International Conference on Data Mining*, Society for Industrial and Applied Mathematics. pp. 473–484. doi:10.1137/1.9781611972795.41.
- [31] Muja, M., Lowe, D.G., 2014. Scalable Nearest Neighbor Algorithms for High Dimensional Data. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 36, 2227–2240. doi:10.1109/TPAMI.2014.2321376.
- [32] Mülâyim, M.O., Arcos, J.L., 2008. Understanding Dubious Future Problems, in: *European Conference on Case-Based Reasoning*, Springer, Berlin, Heidelberg. pp. 385–399. doi:10.1007/978-3-540-85502-6_26.

- [33] Mülâyim, M.O., Arcos, J.L., 2018. Perks of Being Lazy: Boosting Retrieval Performance, in: International Conference on Case-Based Reasoning, Springer Verlag. pp. 309–322. doi:10.1007/978-3-030-01081-2_21.
- [34] Reinsel, D., Gantz, J., Rydning, J., 2018. The Digitization of the World - From Edge to Core. IDC White Paper URL: <https://www.seagate.com/files/www-content/our-story/trends/files/idc-seagate-dataage-whitepaper.pdf>. (Last accessed 20 January 2020).
- [35] Ricci, F., Avesani, P., 1995. Learning a local similarity metric for case-based reasoning, in: International Conference on Case-Based Reasoning, Springer Berlin Heidelberg. pp. 301–312. doi:10.1007/3-540-60598-3_27.
- [36] Schaaf, J.W., 1995. "Fish and Sink" An Anytime-Algorithm to Retrieve Adequate Cases, in: International Conference on Case-Based Reasoning, Springer. pp. 538–547. doi:10.1007/3-540-60598-3_50.
- [37] Serrà, J., Arcos, J.L., 2014. An empirical evaluation of similarity measures for time series classification. *Knowledge-Based Systems* 67, 305–314. doi:10.1016/j.knosys.2014.04.035.
- [38] Smyth, B., Cunningham, P., 1996. The Utility Problem Analysed: A Case-Based Reasoning Perspective, in: European Workshop on Advances in Case-Based Reasoning, Springer. pp. 392–399. doi:10.1007/BFb0020625.
- [39] Smyth, B., Keane, M.T., 1995. Remembering to Forget: A Competence-preserving Case Deletion Policy for Case-based Reasoning Systems, in: International Joint Conference on Artificial Intelligence, Citeseer. pp. 377–382.
- [40] Ueno, K., Xi, X., Keogh, E., Lee, D.J., 2006. Anytime classification using the nearest neighbor algorithm with applications to stream mining, in: International Conference on Data Mining, ICDM, IEEE. pp. 623–632. doi:10.1109/ICDM.2006.21.
- [41] 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, 275–309. doi:10.1007/s10618-012-0250-5.
- [42] Wess, S., Althoff, K.D., Derwand, G., 1993. Using k-d trees to Improve the retrieval step in case-based reasoning, in: European Workshop on Case-Based Reasoning, Springer. pp. 167–181. doi:10.1007/3-540-58330-0_85.
- [43] Wilson, D.C., Leake, D.B., 2001. Maintaining CaseBased Reasoners: Dimensions and Directions. *Computational Intelligence* 17, 196–213. doi:10.1111/0824-7935.00140.
- [44] Woodbridge, J., Mortazavi, B., Bui, A.A., Sarrafzadeh, M., 2016. Improving biomedical signal search results in big data case-based reasoning environments. *Pervasive and Mobile Computing* 28, 69–80. doi:10.1016/j.pmcj.2015.09.006.
- [45] Xu, W., Miranker, D., Mao, R., Ramakrishnan, S., 2008. Anytime K-Nearest Neighbor Search for Database Applications, in: International Conference on Data Engineering Workshop, IEEE. pp. 426–435. doi:10.1109/ICDEW.2008.4498354.
- [46] Yianilos, P.N., 1993. Data Structures and Algorithms for Nearest Neighbor Search in General Metric Spaces, in: SODA, ACM-SIAM Symposium on Discrete Algorithms, Society for Industrial and Applied Mathematics, Philadelphia, PA, USA. pp. 311–321.
- [47] Zilberstein, S., 1993. Operational Rationality Through Compilation of Anytime Algorithms. PhD thesis. University of California at Berkeley.
- [48] Zilberstein, S., 1996. Using Anytime Algorithms in Intelligent Systems. *AI Magazine* 17, 73–83. doi:10.1609/aimag.v17i3.1232.