This project has received funding from SMALL GRANT SCHEME Call under grant agreement NOR/SGS/BIPOLAR/0239/2020-00.

Bipolar disorder prediction with sensor-based semi-supervised learning

D3.3 – Interval and fuzzy approach to affective states modeling

Deliverable No.	D3.3	**Due date**	30-Dec-2023
Type	Report	**Dissemination Level**	Public
Version	1.0	**WP**	WP3
Description	Modeling uncertainty about sensor data and psychiatric labeling of affective states		

Systems Research Institute Polish Academy of Sciences, Newelska 6, 01-447 Warsaw, Poland

Authors

Name	Email
Katarzyna Kaczmarek-Majer	k.kaczmarek@ibspan.waw.pl
Olga Kamińska	o.kaminska@ibspan.waw.pl
Kamil Kmita	kmita@ibspan.waw.pl
Jakub Małecki	jakub.malecki@ibspan.waw.pl

History

Date	Version	Change
12-Dec-2023	0.1	First version of the document
30-Dec-2023	0.2	Version for internal review
22-Jan-2024	1.0	Version ready for submission

Executive summary

This deliverable is output of Task 3.1 and Task 3.2 activities dedicated to modeling uncertainty via interval and fuzzy numbers about data samples and its labeling enabling joint processing of sensor and psychiatric data. The primary goal of this report is illustrating the performance and interpretability of the considered fuzzy approaches for modeling affective states and the related imprecision. The proposed methods address the need of the considered practical context of monitoring of bipolar disorder states. So far in the state of the art, mainly crisp methods have been applied, e.g., for the establishment of the goundtruth period. However, due to the multiple sources of uncertainty in the consider mental health monitoring context, crisp methods often result ineffective. Therefore, we propose several methodologies based on fuzzy set theory that enable to directly incorporate imprecision into the preprocessing and modeling.

The methodological advancements of this deliverable will be next applied in multiple practical evaluations and scenarios for the considered use case of sensor and psychiatric data D1.3. ("Final technical evaluation") and D1.5 ("Final use case evaluation") in Tasks within work package 1. Further comparative analysis (including also crisp methods) will illustrate important practical benefits and limitations of the proposed approaches.

This deliverable is also closely related to the software component developed under WP4 ("Integrated BIPOLAR software package"). It needs to be noted that all theoretical approaches described in this deliverable have been thoroughly tested, carefully implemented and are available for public use as an open-source software components.

Contents

List of acronyms

Acronym	Explanation
BIPOLAR	Bipolar disorder prediction with sensor-based semi-supervised learning project
BD	Bipolar disorder
ADP	Acoustic Data Pilot
LDP	Locomotor Data Pilot
SSFC	Semi-supervised fuzzy clustering
SSFCM	Semi-supervised fuzzy c-means algorithm
SSL	Semi-supervised learning
SSMC-FCM	Algorithm with Multiple Fuzzification Coefficients
DISSFCM	Dynamic Incremental Semi-Supervised Fuzzy C-Means
S3FCM	Safe Semi-Supervised Fuzzy C-Means
CPR	Confidence Path Regularization
RPCM	Repulsive Possibilistic C-Means
PFCM	Possibilistic Fuzzy C-Means
SSPFCM	Semi-Supervised Possibilistic Fuzzy C-Means
PCM	Possibilistic C-Means algorithm
PRISMA	Preferred Reporting Items for Systematic Reviews and Meta-Analyses

1 Introduction

The aim of this report is to summarize the main findings from Tasks 3.1 and 3.2 conducted between months 1 and 24. These tasks were focused on incorporating imprecision about psychiatric labeling through fuzzy modeling in the context of sensor-based bipolar disorder (BD) monitoring.

The primary achievements of this deliverable encompass innovative approaches and considerations within semi-supervised learning methods, specifically tailored for addressing uncertainty in fuzzy clustering. The report also delves into theoretical considerations regarding effective and explainable steps in the realm of semi-supervised fuzzy clustering.

2 Related Work

The challenge of developing models for monitoring bipolar disorder typically involves the remote collection of patient data and stationary psychiatric assessments. We designate the former as `sensors` data, encompassing elements such as self-reported questionnaires [1], data from wearable devices [2], or acoustic features from phone calls collected by smartphones [3]. The latter, i.e., data from psychiatrists, is referred to as `psychiatric` data. Further on, without loss of generality, we confine `sensors` data to summaries of p acoustic features gathered from phone calls.

The `sensors` data is unsupervised. We observe phone call summaries $x_j \in \mathbb{R}^p$ for $j = 1, \ldots, N$, lacking information about the health status of patients during the time these calls transpired.

The `psychiatric` dataset contains information about health status. Labels are encoded as a categorical response variable y denoting the class to which an observation belongs. For instance, a psychiatrist evaluates a patient and categorizes their status using CGI-BD classification, yielding $y \in \{\text{mania}, \text{euthymia}, \text{mixed}, \text{depression}\}$.

A form of partial supervision we do not consider is constraint. In such cases, information is provided of the nature "observations x_0 and x_1 belong to the same class" or, conversely, "observations x_0 and x_1 must not belong to the same class." The actual label of the considered class is not necessarily provided. We further refer to "partial supervision in the form of labels" as "partial supervision."

Typically, the two data sources are merged to extrapolate labels y onto the unsupervised sensor summaries x_j based on common factors such as patient identifier and temporal proximity. To measure temporal proximity, the so-called ground truth period is assumed [4]. For example, summaries of phone calls within a time window spanning from 7 days before the visit up to 2 days after the visit are annotated with the label y obtained during the visit; the remaining data outside of these windows remains unsupervised. Consequently, our problem is categorized as semi-supervised learning (SSL). A comprehensive introduction and overview of SSL methods are provided by [5]. Semi-supervised learning is often characterized as a setting between (completely) unsupervised and (completely) supervised learning. The key feature of SSL is possessing additional knowledge about a subset of M observations out of all available N observations ($M < N$), as is the case in monitoring BD.

In this context, natural questions arise regarding several sources of uncertainty. Firstly, **extrapolation uncertainty** considers temporal closeness. Sensor data collected 7 days before the visit is more questionable in terms of belonging to the same hidden disease phase y as the data collected on the day of the visit. This uncertainty, found to have been neglected in the literature, is discussed by [6] (refer to Section 3.1 for more details).

Secondly, the annotation process related to the ground truth period is of an indirect nature. Psychiatrists do not listen to the actual phone calls due to privacy restrictions. The labels from the `psychiatric` dataset are extrapolated based on the assumption that voice features collected during a certain disease phase share common characteristics and correlate with the response variable y. This assumption may fail, leading to what we term **label uncertainty**.

It's important to note that psychiatrists deliver a single label y from each stationary assessment, and we do not question the validity of the label (e.g., if the psychiatrist provides $y =$ *euthymia*, we do not assert that the correct disease phase is $y =$ *mania*). However, the label uncertainty problem is typically perceived as an issue of erroneous labels. Most methods, such as [7] and [8], as well as literature surveys like [9–11], interpret uncertain labels as corrupted by noise, with the objective of reducing this noise. Recent studies have highlighted the detrimental effects of training deep learning models with noisy annotations [12]. However, we argue that in the context of BD monitoring, one may wish to modify the impact of a given y on the outcomes of the learning procedure (either decrease or increase), without necessarily considering that a different class should have been annotated. This modification of the impact of partial supervision should ideally be data-driven. We proposed an idea for such a procedure in [6], which we termed Confidence Path Regularization.

Semi-Supervised Fuzzy Clustering (SSFC) is a class of semi-supervised models capable of handling the aforementioned uncertainties in a clear and explainable way. SSFC is a cross-functional topic, combining knowledge from various fields: fuzzy logic, cluster analysis, and semi-supervised learning. Researchers contributing to SSFC have diverse backgrounds and study different aspects of the semi-supervised fuzzy clustering problem. However, all various SSFC models, including [6, 13–32], share the same concept of a soft assignment of each jth observation to the kth cluster, $k = 1, \dots, c$, rooted in fuzzy set theory [33]. The implementation of the membership degree $u_{jk} \in [0, 1]$ for the soft assignment enables the analysis of another level of uncertainty - the **data uncertainty**, allowing the modeling of so-called "hybrids" - observations belonging to more than one class (cluster). In contrast, "hard" clustering methods like K-Means or "hard" SSL methods model each observation as belonging unambiguously to a single class only. Therefore, fuzzy sets and the Semi-Supervised Fuzzy Clustering models, built on the concept of fuzzy sets, are the correct tools for modeling the complex uncertainties present in the problem of remote monitoring of bipolar disorder.

3 Handling uncertainty in fuzzy clustering

In their classification of Semi-Supervised Fuzzy Clustering (SSFC) models, [34] categorize them into three groups: (i) distance-based, (ii) constraints-based, and (iii) hybrids. The classification criterion is based on the type of partial supervision and the method used to incorporate it. This discussion focuses specifically on distance-based models, which formalize the concept of similarity between observations and clusters using distance metrics.

In the context of distance-based SSFC models, observations (patterns) are represented by p-vectors denoted as $x_j \in \mathbb{R}^p$, and clusters are represented by p-vectors called prototypes denoted as $v_k \in \mathbb{R}^p$. The objective of SSFC is to formulate an objective function that quantifies distances between observations and patterns, adjusted by memberships. Subsequently, a minimization problem is posed, and an algorithm is proposed to minimize the objective function.

A classical example of a distance-based SSFC model is the Semi-Supervised Fuzzy C-Means proposed by [35], defined by the following objective function:

$$J_{\text{SSFCM}} = \sum_{k=1}^{c} \sum_{j=1}^{N} u_{jk}^2 d_{jk}^2 + \alpha \sum_{k=1}^{c} \sum_{j=1}^{N} b_j (u_{jk} - f_{jk})^2 d_{jk}^2, \tag{1}$$

where: $U = [u_{jk}]_{N \times c}$ is a memberships matrix, $V = [\mathbf{v}_k]_{c \times p}$ is a prototypes matrix, $X = [\mathbf{x}_j]_{N \times p}$ is a features matrix, d_{jk} is the Euclidean distance between an observation and a cluster prototype, $F = [f_{jk}]$ is a matrix introducing partial supervision that contains assumed membership values, $b_j \in \{0, 1\}$ is an indicator variable equal to 1 iff $\mathbf{x}_j$ is labeled, and $\alpha \geq 0$ is a scalar weighting the proportional contribution of partial supervision. Note that $f_{jk} \in \{0, 1\}$.

The objective function J_{SSFCM} can be seen as a penalization adjustment, pushing the membership u_{jk} towards a value of 1 for the supervised observation. Effectively, supervised observations are encouraged to belong to the cluster associated with their class, but they are not necessarily forced to. Note that a supervised observation may be classified into a different cluster than the one associated with its class, allowing for **data uncertainty** identification. This **data uncertainty** enables the identification of observations that deviate from the assumption that `sensors` data correlate with the given label obtained from the `psychiatric` data.

An idea proposed in [35] that we will find very useful in addressing label and extrapolation uncertainty is to incorporate a so-called confidence factor $\text{conf}_j \in [0, 1]$ into the model, resulting in:

$$J_{\text{SSFCM}} = \sum_{k=1}^{c}\sum_{j=1}^{N} u_{jk}^2 d_{jk}^2 + \alpha \sum_{k=1}^{c}\sum_{j=1}^{N} \text{conf}_j b_j (u_{jk} - f_{jk})^2 d_{jk}^2. \tag{2}$$

The confidence factor conf_j is described by [35, p. 790] as follows:

"The higher the confidence level, the more significant the contribution of the corresponding pattern to the objective function."

However, this confidence factor, in its original form, needs to be provided a priori by the user of the model. It is noteworthy that the confidence factor conf_j is not an independent factor, as it simply modifies the scaling factor α. Consequently, observation-specific scaling factors $\alpha_j = \alpha \cdot \text{conf}_j$ can be introduced, and Equation 2 can be rewritten accordingly.

$$\sum_{k=1}^{c}\sum_{j=1}^{N} u_{jk}^2 d_{jk}^2 + \sum_{k=1}^{c}\sum_{j=1}^{N} \alpha_j (u_{jk} - f_{jk})^2 d_{jk}^2. \tag{3}$$

3.1 Confidence Path Regularization

In our recent work [6], we tackled the challenges of extrapolation and label uncertainties inherent in the context of monitoring Bipolar Disorder (BD). This research introduced novel methodologies to improve the accuracy and reliability of the ground truth period and proposed an algorithm named Confidence Path Regularization to estimate adjusted confidence factors for observations in the monitoring process.

Firstly, we identified limitations in the conventional binary approach used to establish the ground truth period. The binary nature of labeling, where phone call summaries recorded within a specific window were treated uniformly, posed challenges in reflecting the temporal aspects of data collection. To overcome this, we suggested employing continuous functions to gradually decrease the confidence factor for observations recorded further from the visit. Figure 1 illustrates the inadequacy of binary annotation and presents alternative strategies, including step and Gaussian functions, to address this limitation.

Secondly, we introduced the Confidence Path Regularization algorithm to compute data-driven, adjusted confidence factors $\text{conf}_j^\star$. This algorithm is grounded in the regularization assumption that highly certain supervised observations

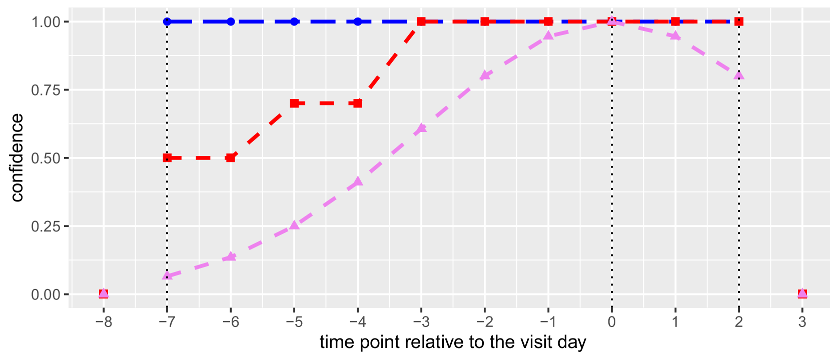

Figure 1: Binary (blue dots), step (red squares), and Gaussian (violet triangles) ground truth period extrapolation strategies. Each strategy implements confidence as a function of time.

should consistently receive high degrees of membership to the supervised class by the Semi-Supervised Fuzzy C-Means (SSFCM) method across varying confidence factor values [6].

The algorithm involves key parameters such as the number of regularization passes (R), the strength of each regularization ($\text{reg}_r \in [0, 1]$), and the weight assigned to the outcome of each regularization pass ($w_r \in \mathbb{N}$). During each pass ($r = 1, \ldots, R$), the confidence factor conf_i for each observation is scaled by the regularization factor reg_r before fitting the model. The supervised class membership values obtained from the R models are then weighted by scalars w_r. The set of regularization factors $\{\text{reg}_r\}_{r=1,\ldots,R}$ forms the confidence regularization path, and, in conjunction with the weights w_r, they constitute the parameters of the proposed method.

The adjusted confidence factor, denoted as $\text{conf}_i^\star$, is obtained through the normalization process expressed in Equation 4, where $s(i)$ is an index of the column in the F matrix for the i-th observation that contains the ground truth-based label.

$$\text{conf}_i^\star = \frac{1}{\sum_{r=1}^{R} w_r} \cdot \sum_{r=1}^{R} u_{i,s(i)}^r w_r, \tag{4}$$

Table 1 provides exemplary data required to calculate the adjusted confidence factor $\text{conf}_{i_0}^\star$ for a given observation i_0 in a procedure composed of $R = 2$ steps.

Table 1: Exemplary values of CPR procedure.

r	conf_{i_0}	reg_r	$u_{i_0,s(i_0)}^r$	w_r
1	1.0	0.1	0.7	10
2	1.0	0.5	0.8	2

Below we present the calculations that involve normalizing the weighted sum of supervised class membership values, resulting in the adjusted confidence factor $\text{conf}_{i_0}^\star \approx 0.72$ for the given observation i_0. This value indicates a high degree of certainty for the supervised observation.

$$
\begin{aligned}
\text{conf}^{\star}_{i_0} &= \frac{1}{\sum_{r=1}^{2} w_r} \cdot \sum_{r=1}^{2} u^r_{i_0,s(i_0)} w_r = \frac{1}{w_1 + w_2} \cdot \left(u^1_{i_0,s(i_0)} \cdot w_1 + u^2_{i_0,s(i_0)} \cdot w_2\right) \\
&= \frac{1}{10+2} \cdot \left(0.7 \cdot 10 + 0.8 \cdot 2\right) = \frac{1}{12} \cdot \left(7 + 1.6\right) = \frac{8.6}{12} \approx 0.72.
\end{aligned}
$$

In conclusion, the adjusted confidence factors, derived through the Confidence Path Regularization algorithm, serve as valuable outcomes for distinguishing observations with varying levels of certainty. These adjusted confidence factors can be utilized as final results for analyses or as data-driven inputs for the Semi-Supervised Fuzzy C-Means model in future applications.

3.2 Explainable Impact of Partial Supervision

As discussed in Section 3, the confidence factor conf_j in the context of the Semi-Supervised Fuzzy C-Means (SSFCM) model is not independent; it modifies the scaling factor α, resulting in observation-specific values α_j. Given that α is the primary hyperparameter regulating the impact of partial supervision in the SSFCM model, it becomes crucial for uncertainty-handling procedures, such as the Confidence Path Regularization (CPR) approach, to gain a comprehensive understanding of its role in the modeling process.

Surprisingly, we identified a gap in the literature concerning a unified and mathematically rigorous framework for setting the value of α. In [36], we addressed this gap by introducing an explainability framework and providing novel insights into the role of α in both the Semi-Supervised Fuzzy C-Means and Semi-Supervised Possibilistic C-Means models.

Our argument challenges the implicit linear treatment of the impact of partial supervision as a function of α and establishes this impact as a non-linear function of α. This non-linear quantification holds significant implications for the Confidence Path Regularization approach, as the weights w_r should appropriately reflect this non-linear impact.

The proposed explainability framework defines an explanation as a description that satisfies three criteria:

(C1) interpretability,

(C2) completeness,

(C3) quantification.

This framework distinguishes between explanations and interpretations, where an interpretation is any description that satisfies criterion (C1) and at least one more criterion (C1 or C2), but not all three criteria. The superiority of explanations over interpretations is emphasized in our framework.

We adopt criteria (C1) and (C2) from Gilpin et al. [37], and introduce criterion (C3) as an additional requirement specific to the scaling factor α: we aim to express the impact of partial supervision as a function IPS(α).

The established interpretations of α in the literature were similar, adhering to the core meaning provided by Pedrycz and Waletzky in [35]. Therefore, we consider the interpretation from [35] as canonical, formulating Interpretation 1.

Interpretation 1 (Pedrycz97) *The role of the scaling factor α is to maintain a balance between the supervised and unsupervised components within the optimization mechanism.*

However, our investigation reveals that Interpretation 1 may implicitly assume $\text{IPS}_{P97}(\alpha) = \alpha$, as "balance" can be understood in terms of "direct proportionality". We, therefore, delve into the relationship between α and the key outcomes of the SSFCM model, namely the estimated memberships $\hat{U}$. The formula for the estimated membership of the jth observation to the kth cluster $\hat{u}_{jk}$ is presented in Equation 4.

$$\hat{u}_{jk} = \frac{1 + \alpha\big(b_j - b_j \sum_{g=1}^{c} f_{jg}\big)}{1 + \alpha b_j} \frac{1}{\sum_{g=1}^{c} d_{jk}^2 / d_{jg}^2} + \frac{\alpha b_j}{1 + \alpha b_j} f_{jk}. \tag{4}$$

We focus on the second component of Equation 4 and propose a novel explanation:

Explanation 1 (IPS in SSFCMeans) *The scaling factor α quantifies the impact of partial supervision as IPS(α) = $\frac{\alpha}{1+\alpha}$, establishing an Absolute Lower Bound for the membership of a supervised observation to the supervised cluster, ensuring $u_{i,s(i)} >$ IPS(α).*

Figure 2 visualizes the IPS(α)=$\frac{\alpha}{1+\alpha}$ function along with its derivative $\frac{\partial \text{SSFCM}}{\partial \alpha}$. The proposed explainability framework in [36] offers substantial benefits: clear criteria distinguish interpretations from explanations, and the functional form of the impact of partial supervision leads to new insights into the mechanism of partial supervision, enabling enhanced procedures. For CPR, we recommend improvements in the weighing mechanism: weights w_r should consider the $\text{IPS}(\alpha) = \frac{\alpha}{1+\alpha}$, rather than being set to the inverse of the regularization factor reg_r.

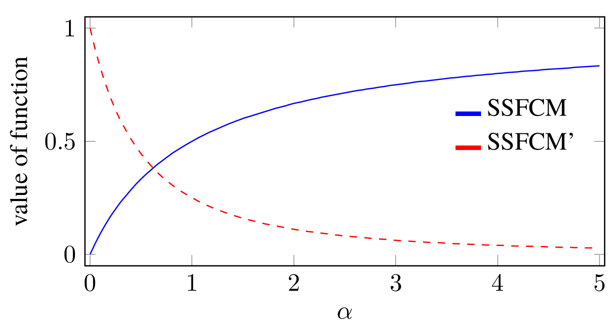

Figure 2: The impact of partial supervision IPS(α) for $\alpha \in [0, 5]$ for both SSFCMeans (blue solid line) and its derivative SSFCM' (red dashed line).

3.3 Classification Error in Semi-Supervised Fuzzy C-Means

In accordance with Chapelle et al.'s comprehensive exploration of Semi-Supervised Learning (SSL) [5, p. 16], wherein SSL is categorized as a supervised problem aiming to minimize classification error, our research focuses on the utilization of the Semi-Supervised Fuzzy C-Means (SSFCMeans) model as a classifier. The primary objective is to induce a rule $\pi : x_j \mapsto y_k$ for classifying observations, guided by the cluster assumption that points within the same cluster are likely to belong to the same class [5, p. 5].

The output of the SSFCMeans model for the ith observation is represented as a tuple $\langle \hat{u}_{i1}, \ldots, \hat{u}_{ic} \rangle$. To determine the predicted class, a defuzzification process is required. A straightforward approach involves selecting the class associated with the cluster having the highest membership value. This decision is formalized by the decision rule π defined as:

$$\hat{\pi} : i \mapsto k^\star = \arg\max_k \hat{u}_{ik}. \tag{5}$$

In the context of classification, an analysis of the relationship between training error and test error, as well as the consideration of overfitting, becomes crucial. The induction of the classifier $\hat{\pi} : x_j \mapsto y_j$ from the training data $\mathcal{D}$ may lead to a classifier that performs exceptionally well on $\mathcal{D}$ but lacks generalization capability, resulting in decreased performance on independent test data $\mathcal{T}$. Therefore, a comprehensive evaluation of both training and test errors is necessary.

However, as demonstrated in our recent work [38], the overfitting analysis for the Semi-Supervised Fuzzy C-Means model is atypical due to the nature of the training error, which is based on the formula for estimated membership $\hat{u}_{i,s(i)}$ of supervised observations to the supervised cluster (see Eq. 4). Importantly, we establish a guaranteed lower bound as expressed in Equation 6:

$$\hat{u}_{i,s(i)} \geq \frac{\alpha}{1+\alpha}. \tag{6}$$

This lower bound, denoted as $\frac{\alpha}{1+\alpha}$, is termed the Absolute Lower Bound to emphasize its significance.

For any value of $\alpha \geq 1$, we can ensure that the cluster indexed by $s(i)$ will be the argument maximum, as the associated membership $\hat{u}_{i,s(i)} \geq 0.5$. This leads to a 0% training error irrespective of the chosen performance metric for classification error assessment (e.g., precision, recall, or F1 score). Such scenarios challenge traditional overfitting analyses, as the classifier not only fits excessively to the local characteristics of the training data $\mathcal{D}$ but is also deterministically assured to achieve artificially outstanding performance, regardless of the data evidence.

To address this issue, we propose considering a *margin of making the training error* $\epsilon \in [0, 0.5]$ during the selection of the optimal value of α, taking into account both training and test classification performance. In terms of the scaling factor α, we advocate considering α such that

$$\frac{\alpha}{1+\alpha} = 0.5 - \epsilon. \tag{7}$$

It is noteworthy that introducing a margin for making the training error ϵ may still result in 0% training error. However, in such cases, a traditional overfitting analysis becomes meaningful because the performance would no longer be deterministic.

In [38], we provide examples of the analysis, including (A) a curated simulation scenario, and (B) real-life data concerning the monitoring of bipolar disorder.

3.4 Time ordering

Until now, we have only assumed that $x_j \in \mathbb{R}^p$; however, it's common for observations to have a temporal order, given that we know the start time and end time of each phone call. For specific details, please consult Table 2 including exemplary data; calls 2 and 3 occurred within hours on the same day, but the first call took place a couple of days earlier.

Table 2: Exemplary ordered data.

Patient ID	Call ID	Data	Start	End
100	1	$x_1 \in \mathbb{R}^4$	15 October 2023, `15:33:01`	15 October 2023, `15:38:12`
100	2	$x_2 \in \mathbb{R}^4$	19 October 2023, `08:09:01`	19 October 2023, `08:30:52`
100	3	$x_3 \in \mathbb{R}^4$	19 October 2023, `13:14:59`	19 October 2023, `13:20:44`

In general, this additional knowledge about the temporal domain can be leveraged in various ways, extending beyond Semi-Supervised Fuzzy Clustering. Connected fields are named, for example, "data streaming," "online learning," and more.

One approach to exploit this information was proposed in [39]. The authors build on the clustering approach and view the challenge of streaming clustering algorithms as an incremental classification problem. They assume a continuous setting, where data can arrive at any time, and the model should integrate this new data to reconfigure itself. The novel observation needs to be *classified* into one of the existing clusters, or a decision to form a new cluster must be made.

Another strategy for exploiting the temporal structure of the dataset was introduced in [40]. There, the temporal domain is utilized to divide the dataset X of cardinality N into several batches (data chunks) to identify patterns changing over time at the batch level rather than individual observations. In [40], a method called Dynamic Incremental Semi-Supervised Fuzzy C-Means (DISSFCM) is proposed, wrapping the core SSFCM model and applying it at the batch level, appropriately processing and summarizing the results from fitting the model to each batch.

The preliminary idea proposed in this report is of a different nature. We suggest using the temporal order as a form of external supervision. Indeed, the partial supervision we discussed relies on using labels y_i, encoded in the form of f_{ik}, to push the respective supervised memberships $u_{i,s(i)}$ towards the value 1. Similarly, for a given observation i_0, one can adjust the entire distribution of memberships $\{u_{i_0,k}\}_k$ to be close to the distribution of memberships for a preceding observation $\{u_{i_0-1,k}\}_k$. This preliminary idea, however, requires further investigation into its computational and algorithmic aspects. Note that it might be combined with regular partial supervision y_i.

4 Software component `ssfclust`

We developed an R library, `ssfclust` (which stands for semi-supervised fuzzy clustering), to implement two core Semi-Supervised Fuzzy Clustering models, namely:

(i) Semi-Supervised Fuzzy C-Means,

(ii) Semi-Supervised Possibilistic C-Means.

The optimization algorithm and formulas' details align with the description of both models presented in [36]. The open-source library is freely available at `https://github.com/ITPsychiatry/ssfclust` with the appropriate license.

To the best of our knowledge, `ssfclust` represents the first comprehensive open-source implementation of these models in the R language in the form of a downloadable library. While regular packages for unsupervised fuzzy clustering, such as `ppclust` [41] or `fclust` [42], exist, they do not cover the semi-supervised case.

Moreover, the library is entirely written in R, without the need for compiling C++ code, a common practice for improving code efficiency. This approach enhances code readability for practitioners and facilitates easy extension of these models without the necessity to rewrite crucial parts of the algorithms shared among all derivative models.

References

[1] J. Bardram, M. Frost, K. Szántó, M. Faurholt-Jepsen, M. Vinberg, and L. Kessing, "Designing mobile health technology for bipolar disorder: A field trial of the monarca system," 04 2013, pp. 2627–2636.

[2] G. Valenza, C. Gentili, A. Lanatà, and E. P. Scilingo, "Mood recognition in bipolar patients through the psyche platform: Preliminary evaluations and perspectives," Artificial Intelligence in Medicine, vol. 57, no. 1, pp. 49–58, 2013. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0933365712001492

[3] M. Dominiak, K. Kaczmarek-Majer, A. Z. Antosik-Wójcińska, K. R. Opara, A. Olwert, W. Radziszewska, O. Hryniewicz, Ł. Święcicki, M. Wojnar, and P. Mierzejewski, "Behavioral and self-reported data collected from smartphones for the assessment of depressive and manic symptoms in patients with bipolar disorder: Prospective observational study," J Med Internet Res, vol. 24, no. 1, p. e28647, Jan 2022. [Online]. Available: https://www.jmir.org/2022/1/e28647

[4] A. Grünerbl, A. Muaremi, V. Osmani, G. Bahle, S. Öhler, G. Tröster, O. Mayora, C. Haring, and P. Lukowicz, "Smartphone-based recognition of states and state changes in bipolar disorder patients," IEEE Journal of Biomedical and Health Informatics, vol. 19, no. 1, pp. 140–148, 2015.

[5] O. Chapelle, B. Schölkopf, and A. Zien, Eds., Semi-Supervised Learning, ser. Adaptive Computation and Machine Learning. Cambridge, Mass: MIT Press, 2006.

[6] K. Kmita, G. Casalino, G. Castellano, O. Hryniewicz, and K. Kaczmarek-Majer, "Confidence path regularization for handling label uncertainty in semi-supervised learning: use case in bipolar disorder monitoring," in 2022 IEEE International Conference on Fuzzy Systems (FUZZ-IEEE), 2022, pp. 1–8.

[7] D. T. Nguyen, C. K. Mummadi, T. P. N. Ngo, T. H. P. Nguyen, L. Beggel, and T. Brox, "Self: Learning to filter noisy labels with self-ensembling," arXiv preprint arXiv:1910.01842, 2019.

[8] J. Li, R. Socher, and S. C. Hoi, "Dividemix: Learning with noisy labels as semi-supervised learning," arXiv preprint arXiv:2002.07394, 2020.

[9] D. Karimi, H. Dou, S. K. Warfield, and A. Gholipour, "Deep learning with noisy labels: Exploring techniques and remedies in medical image analysis," Medical Image Analysis, vol. 65, p. 101759, 2020.

[10] G. Algan and I. Ulusoy, "Image classification with deep learning in the presence of noisy labels: A survey," Knowledge-Based Systems, vol. 215, p. 106771, 2021.

[11] H. Song, M. Kim, D. Park, Y. Shin, and J.-G. Lee, "Learning from noisy labels with deep neural networks: A survey," arXiv preprint arXiv:2007.08199, 2020.

[12] K. Yi and J. Wu, "Probabilistic end-to-end noise correction for learning with noisy labels," in Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 2019, pp. 7017–7025.

[13] C. Stutz and T. Runkler, “Classification and prediction of road traffic using application-specific fuzzy clustering,” IEEE Transactions on Fuzzy Systems, vol. 10, no. 3, pp. 297–308, Jun. 2002.

[14] Chunfang Li, Lianzhong Liu, and Wenli Jiang, “Objective function of semi-supervised Fuzzy C-Means clustering algorithm,” in 2008 6th IEEE International Conference on Industrial Informatics. Daejeon, South Korea: IEEE, Jul. 2008, pp. 737–742.

[15] E. Yasunori, H. Yukihiro, Y. Makito, and M. Sadaaki, “On semi-supervised fuzzy c-means clustering,” in 2009 IEEE International Conference on Fuzzy Systems. Jeju Island, South Korea: IEEE, Aug. 2009, pp. 1119–1124.

[16] A. Bouchachia and W. Pedrycz, “Enhancement of fuzzy clustering by mechanisms of partial supervision,” Fuzzy Sets and Systems, vol. 157, no. 13, pp. 1733–1759, Jul. 2006.

[17] ——, “A Semi-supervised Clustering Algorithm for Data Exploration,” in Fuzzy Sets and Systems — IFSA 2003, J. G. Carbonell, J. Siekmann, T. Bilgiç, B. De Baets, and O. Kaynak, Eds. Berlin, Heidelberg: Springer Berlin Heidelberg, 2003, vol. 2715, pp. 328–337.

[18] J. Gao, P.-N. Tan, and H. Cheng, “Semi-Supervised Clustering with Partial Background Information,” in Proceedings of the 2006 SIAM International Conference on Data Mining. Society for Industrial and Applied Mathematics, Apr. 2006, pp. 489–493.

[19] J. Xu, G. Feng, T. Zhao, X. Sun, and M. Zhu, “Remote sensing image classification based on semi-supervised adaptive interval type-2 fuzzy c-means algorithm,” Computers & Geosciences, vol. 131, pp. 132–143, Oct. 2019.

[20] L. Liu and X.-J. Wu, “Semi-Supervised Possibilistic Fuzzy c-Means Clustering Algorithm on Maximized Central Distance,” in Proceedings of the 2nd International Conference on Computer Science and Electronics Engineering (ICCSEE 2013). China: Atlantis Press, 2013.

[21] F. Salehi, M. R. Keyvanpour, and A. Sharifi, “SMKFC-ER: Semi-supervised multiple kernel fuzzy clustering based on entropy and relative entropy,” Information Sciences, vol. 547, pp. 667–688, Feb. 2021.

[22] G. Casalino, M. Dominiak, F. Galetta, and K. Kaczmarek-Majer, “Incremental Semi-Supervised Fuzzy C-Means for Bipolar Disorder Episode Prediction,” in 2020 IEEE Conference on Evolving and Adaptive Intelligent Systems (EAIS). Bari, Italy: IEEE, May 2020, pp. 1–8.

[23] G. Casalino, G. Castellano, F. Galetta, and K. Kaczmarek-Majer, “Dynamic Incremental Semi-supervised Fuzzy Clustering for Bipolar Disorder Episode Prediction,” in Discovery Science, A. Appice, G. Tsoumakas, Y. Manolopoulos, and S. Matwin, Eds. Cham: Springer International Publishing, 2020, vol. 12323, pp. 79–93.

[24] G. Casalino, G. Castellano, and C. Mencar, “Data Stream Classification by Dynamic Incremental Semi-Supervised Fuzzy Clustering,” International Journal on Artificial Intelligence Tools, vol. 28, no. 08, p. 1960009, Dec. 2019.

[25] K. Kaczmarek-Majer, G. Casalino, G. Castellano, O. Hryniewicz, and M. Dominiak, “Explaining smartphone-based acoustic data in bipolar disorder: Semi-supervised fuzzy clustering and relative linguistic summaries,” Information Sciences, vol. 588, pp. 174–195, Apr. 2022.

[26] H. Gan, Z. Yang, and R. Zhou, "Adaptive safety-aware semi-supervised clustering," Expert Systems with Applications, vol. 212, p. 118751, Feb. 2023.

[27] P. T. Huan, P. H. Thong, T. M. Tuan, D. T. Hop, V. D. Thai, N. H. Minh, N. L. Giang, and L. H. Son, "TS3FCM: Trusted safe semi-supervised fuzzy clustering method for data partition with high confidence," Multimedia Tools and Applications, vol. 81, no. 9, pp. 12 567–12 598, Apr. 2022.

[28] L. Guo, H. Gan, S. Xia, X. Xu, and T. Zhou, "Joint exploring of risky labeled and unlabeled samples for safe semi-supervised clustering," Expert Systems with Applications, vol. 176, p. 114796, Aug. 2021.

[29] H. Gan, Y. Fan, Z. Luo, R. Huang, and Z. Yang, "Confidence-weighted safe semi-supervised clustering," Engineering Applications of Artificial Intelligence, vol. 81, pp. 107–116, May 2019.

[30] H. Gan, Y. Fan, Z. Luo, and Q. Zhang, "Local homogeneous consistent safe semi-supervised clustering," Expert Systems with Applications, vol. 97, pp. 384–393, May 2018.

[31] V. Antoine, J. A. Guerrero, T. Boone, and G. Romero, "Possibilistic clustering with seeds," in 2018 IEEE International Conference on Fuzzy Systems (FUZZ-IEEE). Rio de Janeiro: IEEE, Jul. 2018, pp. 1–7.

[32] V. Antoine, J. A. Guerrero, and G. Romero, "Possibilistic fuzzy c-means with partial supervision," Fuzzy Sets and Systems, vol. 449, pp. 162–186, Nov. 2022.

[33] L. Zadeh, "Fuzzy sets," Information and Control, vol. 8, no. 3, pp. 338–353, Jun. 1965.

[34] D. T. C. Lai and J. M. Garibaldi, "A comparison of distance-based semi-supervised fuzzy c-means clustering algorithms," in 2011 IEEE International Conference on Fuzzy Systems (FUZZ-IEEE 2011). Taipei, Taiwan: IEEE, Jun. 2011, pp. 1580–1586.

[35] W. Pedrycz and J. Waletzky, "Fuzzy clustering with partial supervision," IEEE Transactions on Systems, Man, and Cybernetics, Part B (Cybernetics), vol. 27, no. 5, pp. 787–795, Oct. 1997.

[36] K. Kmita, K. Kaczmarek-Majer, and O. Hryniewicz, "Explainable impact of partial supervision in semi-supervised fuzzy clustering," revised manuscript submitted for review.

[37] L. H. Gilpin, D. Bau, B. Z. Yuan, A. Bajwa, M. Specter, and L. Kagal, "Explaining Explanations: An Overview of Interpretability of Machine Learning," in 2018 IEEE 5th International Conference on Data Science and Advanced Analytics (DSAA). Turin, Italy: IEEE, Oct. 2018, pp. 80–89.

[38] K. Kmita, K. Kaczmarek-Majer, and O. Hryniewicz, "Classification error in semi-supervised fuzzy c-means," in Fuzzy Logic and Technology, and Aggregation Operators, S. Massanet, S. Montes, D. Ruiz-Aguilera, and M. González-Hidalgo, Eds. Cham: Springer Nature Switzerland, 2023, pp. 725–736.

[39] W. Wu, J. M. Keller, J. Dale, and J. C. Bezdek, "StreamSoNG: A Soft Streaming Classification Approach," IEEE Transactions on Emerging Topics in Computational Intelligence, vol. 6, no. 3, pp. 700–709, Jun. 2022.

[40] G. Casalino, G. Castellano, F. Galetta, and K. Kaczmarek-Majer, "Dynamic Incremental Semi-supervised Fuzzy Clustering for Bipolar Disorder Episode Prediction," in Discovery Science, A. Appice, G. Tsoumakas, Y. Manolopoulos, and S. Matwin, Eds. Cham: Springer International Publishing, 2020, vol. 12323, pp. 79–93.

[41] Z. Cebeci, "Comparison of internal validity indices for fuzzy clustering," Journal of Agricultural Informatics, no. 2, pp. 1–14, 2019.

[42] M. Ferraro, P. Giordani, and A. Serafini, "fclust: An r package for fuzzy clustering," The R Journal, vol. 11, 2019. [Online]. Available: https://journal.r-project.org/archive/2019/RJ-2019-017/RJ-2019-017.pdf