

Longitudinal Monitoring of Muscle Weakness in ALS Using Portable sEMG and Data Mining

Abraham Jorge Jiménez Alfaro^{1,2}, Norma Karen Valencia Vázquez², Griselda Cortés Barrera¹, Edgar Corona Organiche¹, Mercedes Flores Flores¹, Irving Alcocer Guillermo Cardiel¹

¹ Tecnológico Nacional de México / Tecnológico de Estudios Superiores de Ecatepec (TESE)/ Laboratorio Nacional CONAHCYT en Inteligencia Artificial y Ciencia de Datos, Ecatepec, Estado de México, México.

² Tecnológico Nacional de México / Tecnológico de Estudios Superiores de Chimalhuacán, Estado de México, México.

E-mails: ajimenez@tese.edu.mx, karenvalencia@teschi.edu.mx, gcortes@tese.edu.mx, ecorona@tese.edu.mx, mflores@tese.edu.mx, ialcocer@tese.edu.mx

Abstract. Amyotrophic Lateral Sclerosis (ALS) is a progressive neurodegenerative disease characterized by the degeneration of upper and lower motor neurons and the gradual loss of muscle function. This article presents an exploratory longitudinal monitoring framework for a single patient with ALS using eight-channel surface electromyography (sEMG) acquired with a Myo Armband and interpretable data-mining procedures. The workflow combines multichannel acquisition, temporal segmentation, time-domain feature extraction (RMS, MAV and WL), normalization, exploratory clustering and a rule-based decision tree inspired by CART. The current implementation is not a data-trained CART model: the thresholds are pilot, auditable criteria derived from the descriptive historical range and normalized reference zones, and they require recalibration in larger labelled cohorts. Because the legacy archive available for this revision does not contain a complete export of session-level and window-level counts or an independent labelled validation set, no population-level classification-performance claim is made. The framework emphasizes within-subject longitudinal comparison, acquisition metadata, quality control and interpretable alerts. The 200 Hz sampling frequency is adequate for the selected time-domain descriptors but constrains fine spectral analysis. Accordingly, the proposed rules should be interpreted as complementary monitoring support and require external clinical or physiotherapeutic validation before diagnostic use.

Keywords: Surface electromyography; Amyotrophic Lateral Sclerosis; longitudinal monitoring; data mining; decision rules; CART.

Article information

Received: Jun 26, 2026

Accepted: Aig 11, 2026

1 Introduction

Fatigue and muscle weakness are phenomena of clinical and technological interest because they reflect functional changes in the contractile capacity of skeletal muscle. In patients with ALS, disease progression leads to motor deterioration and functional dependence; therefore, quantitative longitudinal evidence may complement clinical follow-up. Surface electromyography (sEMG) records muscular electrical activity through electrodes placed on the skin, and its interpretation depends on the acquisition protocol, sensor placement, signal quality and subsequent processing (De Luca, 1997; Hermens et al., 2000; Merletti & Parker, 2004). The present study uses a portable Myo Armband to collect upper-limb sEMG and preserve historical measurements. The analytical contribution is an interpretable, rule-based decision framework informed by decision-tree principles. Decision trees are useful because their if-then paths make explicit which variables and thresholds drive an output (Breiman et al., 1984; Quinlan, 1986). In this manuscript, however, the thresholds are not learned from a labelled training set; they are exploratory pilot rules that must be distinguished from a statistically fitted CART model. The workflow therefore combines: a) multichannel sEMG acquisition; b) temporal feature extraction; c) exploratory pattern analysis; d) rule-based classification inspired by CART; and e) longitudinal interpretation and external clinical review.

2 Materials and Method

2.1 Case Study and Data Acquisition

The case study is conceived as an instrumental, longitudinal and exploratory single-patient study aimed at demonstrating the methodological feasibility of integrating sEMG signals, historical storage and interpretable data mining for functional monitoring. Its purpose is methodological replication rather than statistical generalization (Yin, 2018). The participant was a 57-year-old male patient with a reported diagnosis of Amyotrophic Lateral Sclerosis of three years and six months' progression and reported muscle weakness ranging between 60% and 70%. The available historical material spans two years and focuses on upper-limb signals during a simple motor task. Each session was designed to contain four repetitions of the rest-contraction cycle. The source materials available for this revision do not include a complete machine-readable export of session identifiers or window-level records; therefore, the exact total number of recording sessions, retained repetitions and valid analysis windows cannot be verified retrospectively and are not reported as if known. Under the segmentation specification defined in Section 2.2 (1.0-s windows with 50% overlap), a 15-30 s contraction produces 29-59 candidate windows per repetition, or 116-236 candidate windows per complete four-repetition session before quality exclusions. Future external validation must report the exact session count, valid repetitions, retained windows and their dates directly from the exported database. The single-patient design is appropriate only for within-subject exploratory comparison; it cannot establish population-level ALS thresholds. From a physiological perspective, sEMG represents the spatial and temporal summation of motor-unit action potentials detected at the skin surface, and its amplitude and stability depend on placement, skin impedance, muscle geometry, contraction level, fatigue, movement artefacts and environmental noise (De Luca, 1997; Hermens et al., 2000; Merletti & Parker, 2004). The recording was performed at the patient's home to reduce mobility barriers, but this setting requires strict metadata on forearm position, armband placement, skin preparation, rest/contraction phases, session date and assistance. The Medical Research Council scale is used only as an ordinal clinical reference and does not convert the sEMG output into a diagnosis (Medical Research Council, 1943).

The signals were recorded using the Myo Armband sensor, see Figure 1, which provides eight surface-EMG channels around the forearm and records EMG at 200 Hz. Peer-reviewed descriptions of the device report eight dry sEMG sensors, wireless transmission and a 200 Hz EMG sampling rate (Zhang et al., 2019; Phinyomark et al., 2018). The eight-channel configuration provides a spatial reading around the forearm. Sensors 03 and 04 were treated as priority channels in the legacy analysis because they showed greater relative stability for the recorded task; this selection is exploratory and must be re-evaluated in additional participants. The 200 Hz rate supports low-cost time-domain descriptors such as RMS, MAV and WL, but its lower bandwidth relative to clinical EMG systems restricts fine spectral analysis and motivates a focus on interpretable time-domain features (Phinyomark et al., 2018). Small changes in sensor placement or task execution may materially alter amplitude, so acquisition conditions must be documented before sessions are compared (Hermens et al., 2000; Merletti & Parker, 2004).




Fig. 1. Myo Armband sensor used for sEMG acquisition

Figure 2 describes the conceptual model that summarizes the complete chain through which a patient's motor action is transformed into analytical information. The first stage corresponds to the patient and the functional task: fist closure or palm extension produces changes in muscular electrical activation. In sEMG terms, these changes represent the aggregated activity of motor units detected at the surface; therefore, signal quality depends on the protocol, sensor placement and contraction stability (De Luca, 1997; Hermens et al., 2000).




Fig. 2. Conceptual model of the acquisition, transmission, storage, and analysis workflow

The second stage is represented by the Myo Armband. The device captures eight-channel sEMG from the forearm and transmits the samples wirelessly to a computer. This stage fulfils two methodological functions: multichannel acquisition and preservation of the temporal relationship between samples. Wireless use is practical for home monitoring, but connection stability, packet loss, battery level, device position and session consistency must be checked and recorded (Zhang et al., 2019).

The third stage corresponds to immediate computational processing. The computer receives the signal, organises it by channel and time, and enables it to be segmented into windows associated with resting or contraction phases. Features such as RMS, MAV and WL are calculated from these windows. This step is central because data-mining models do not operate directly on the complete raw signal, but rather on variables that summarize energy, amplitude and temporal variability (Phinyomark et al., 2012).

The fourth stage represents historical storage and consultation. The database icon indicates that each session should retain the date, phase, channel, extracted features and, where possible, clinical or physiotherapy observations. This structure makes it possible to compare sessions and construct longitudinal trajectories. In data-mining terms, ordered storage is the prior condition for applying clustering, training decision trees and auditing classification rules (Han et al., 2012; Witten et al., 2017).

The final stage corresponds to analysis and interpretation. Graphs, CART rules and output tables transform the recorded signal into interpretable evidence for monitoring. At this point, the system may identify windows with lower activation, anomalous

variability or an increase in the D_w index. Nevertheless, the conceptual model should be understood as support for clinical review: the generated rules require validation against specialist assessment and do not replace medical diagnosis. Table 1 summarizes the components, the methodological function within the analysis and the theoretical basis of each stage.

Table 1. Components of the conceptual model for acquisition, transmission, storage, and analysis

Component	Function	Bibliographic support
Patient and motor task	Generate comparable resting and contraction conditions to observe changes in muscle activation.	De Luca (1997); Merletti and Parker (2004)
Myo Armband sensor	Acquire multichannel sEMG signals in a portable and non-invasive manner.	Zhang et al. (2019); Phinyomark et al. (2018); Hermens et al. (2000)
Wireless transmission and computer	Receive the signal, organize it by channel, and segment it into analysis windows.	Han et al. (2012); Witten et al. (2017)
Historical database	Store sessions, dates, channels, and features for longitudinal follow-up.	Han et al. (2012)
Analytics and visualization	Apply clustering, CART trees, and rule interpretation to support clinical decisions.	Breiman et al. (1984); Kotsiantis (2013)

2.2 Sampling and Preprocessing Protocol

The capture session was structured into resting and voluntary-contraction phases and four repetitions per session. Phase 1 records minimal activation with the forearm supported and the palm extended; phase 2 records fist closure. For the reproducible analytical implementation defined in this revision, each channel is processed as a chronological series and segmented into 1.0-s windows (200 samples at 200 Hz) with 50% overlap. The legacy figures are retained in the form in which they were exported; because the historical archive does not document a uniform digital band-pass or notch filter applied to every capture, this revision does not retroactively claim an unverifiable filtering stage. Instead, window-level quality control consists of removal of the local mean (DC offset) before feature calculation, exclusion of windows containing missing or non-finite samples or documented connection failure, and explicit flagging of visibly clipped or corrupted segments. High variability is not automatically discarded because it may reflect either a physiological abrupt event or technical artefact; such windows are routed to a review alert through WL_z,w. This distinction between verified legacy processing and the corrected reproducible protocol prevents unsupported retrospective assumptions.



Fig. 3. Phase 1 of the protocol: extended palm with the sensor placed on the forearm

Phase 2, corresponding to fist closure for 15 to 30 seconds, captures a simple voluntary contraction. With 1.0-s windows and 50% overlap, a 15-s contraction yields 29 candidate windows and a 30-s contraction yields 59 candidate windows before quality exclusions. This phase is compared with the resting baseline under equivalent placement and posture. RMS and MAV summarize activation amplitude, whereas WL contributes a variability descriptor; interpretation is based on their longitudinal change rather than on any single raw sample. A contraction that cannot be completed consistently, or a segment affected by connection failure or missing data, is excluded from feature computation and recorded as a quality-control event.

Phase 3 consists of four repetitions per recorded session and is intended to reduce dependence on a single contraction. The longitudinal archive spans two years, and sessions are intended to be compared by date using the historical database. The exact number and temporal distribution of completed sessions cannot be reconstructed from the manuscript file alone because a complete export of session IDs and dates was not supplied with the legacy material; therefore, these quantities are stated as a current data-provenance limitation rather than estimated. For any subsequent validation release, the database export should

include session date, repetition number, channel, candidate-window count, excluded-window count and final valid-window count. Only sessions acquired under equivalent protocol conditions should be compared longitudinally.

Table 2. Methodological reading and quality control of the experimental protocol

Protocol element	Methodological function	Associated variable or feature	Quality-control criterion
Phase 1: rest	Establishes a baseline and resting noise level.	Baseline RMS, baseline MAV, WL stability.	Verify low activation and absence of relevant artifacts.
Phase 2: closed fist	Records a comparable voluntary contraction.	Active RMS, active MAV, active WL, D_w; 1.0-s windows with 50% overlap.	Exclude windows with missing/non-finite samples or documented connection failure; record all exclusions.
Phase 3: four repetitions	Reduces dependence on a single measurement.	Four repetitions; candidate and valid window counts; dispersion and outlier flags.	Review consistency between cycles and sensors.
Phase 4: historical follow-up	Enables longitudinal analysis by date and session.	Session dates; longitudinal trends in RMS, MAV, WL, D_w and rule outputs.	Compare only sessions with equivalent protocol and complete metadata.

Overall, Table 2 defines the traceability requirements between the motor gesture, the sEMG signal and the subsequent rule-based classification. Changes in armband position, contraction duration, support surface, time of day, prior fatigue or assistance must be retained as metadata. Missing samples are not imputed for feature extraction: any affected window is excluded and the exclusion is counted. Parameters used for normalization (the mean and standard deviation of each feature, and the reference RMS_min and RMS_max used by D_w) must be estimated only from the designated baseline/reference training subset and then applied unchanged to subsequent or validation windows. In future cross-validation of a data-trained CART model, these parameters must be recomputed inside each training fold to avoid information leakage.

2.3 Mathematical Representation of sEMG Signals

Let S be the multichannel record of sEMG signals, where each channel corresponds to one device electrode and each temporal observation represents a relative voltage sample. For computational analysis, the signal is segmented into windows of length N_w , and each window is transformed into a feature vector.

2.3.1 Multichannel representation of the sEMG signal

The expression $S = \{s_e(t) \mid e = 1, \dots, 8; t = 1, \dots, N\}$ defines the complete set of measurements captured by the armband. The subscript e identifies the electrode or channel, while t identifies the sampling instant. As the Myo Armband sensor has eight electrodes, the signal is not analysed as a single time series, but rather as a matrix of eight synchronised channels. This formalisation makes it possible to compare the behaviour of different areas of the forearm and to recognise that some sensors, such as channels 03 and 04, may provide more stable information for the recorded motor task:

$$S = \{S_e(t) \mid |e| = 1, \dots, 8; t = 1, \dots, N\} \quad (1)$$

2.3.2 Feature vector per temporal window

The expression $x_w = [f_1(w), f_2(w), \dots, f_p(w)]^T$ transforms a window of raw signal into a numerical vector. Each $f_j(w)$ represents a feature extracted from window w , for example RMS, MAV or WL. The superscript T indicates that the set of features is organised as a column vector. This conversion is essential because clustering and decision-tree algorithms do not work directly with the entire raw signal, but rather with compact descriptors that summarise amplitude, energy, variability or temporal complexity:

$$\mathbf{x}_w = [f_1(w), f_2(w), \dots, f_p(w)]^T \quad (2)$$

Table 3 summarizes the function of each parameter:

Table 3. Variables of the representation and segmentation formulas

Parameter	Description	Domain
S	Set of recorded sEMG signals.	Multichannel matrix.
s _e (t)	Signal value of electrode e at time t.	Discrete signal value.
e	Electrode or channel index.	1, ..., 8.
t	Sampling-time index.	1, ..., N.
N	Total number of samples captured per channel.	Positive integer.
x _w	Feature vector of window w.	Vector in R ^p .
f _j (w)	j-th feature calculated in window w.	RMS, MAV, WL, or another feature.
p	Number of features used by the model.	Positive integer.

2.4 Feature Extraction and Normalization

Temporal features condense the sEMG activity of each window and reduce the dimensionality of the raw recording. RMS, mean absolute value and waveform length are widely used time-domain descriptors because they capture effective amplitude, average activation and temporal variation (Phinyomark et al., 2012). These descriptors are also suitable for wearable systems with lower sampling rates, although the 200 Hz bandwidth of the Myo Armband limits detailed spectral characterization (Phinyomark et al., 2018).

2.4.1 Root Mean Square, RMS

The RMS_{e,w} is calculated by squaring each sample s_{e,w}(t), averaging those squared values within the window and applying the square root to the result. By squaring the signal, cancellation between positive and negative values is avoided; therefore, RMS functions as a measure of energy or effective activation. In the study, a reduced RMS during the fist-closure phase may indicate lower muscular activation capacity, provided that sensor placement and task execution are comparable across sessions:

$$RMS_{e,w} = \sqrt{\frac{1}{N_w} \sum_{t=1}^{N_w} s_{e,w}(t)^2} \quad (3)$$

2.4.2 Mean Absolute Value, MAV

The MAV_{e,w} obtains the average of the absolute values of the signal within the window. Unlike RMS, it does not emphasise peaks to the same extent because it does not square the samples; therefore, it provides a more direct reading of the average amplitude. In longitudinal monitoring, RMS and MAV should be analysed together: if both decrease during voluntary contraction, the hypothesis of reduced muscular activation is strengthened; if only RMS increases, an isolated peak or spasm may be present:

(4)

$$MAV_{e,w} = \frac{1}{N_w} \sum_{t=1}^{N_w} |s_{e,w}(t)|$$

2.4.3 Waveform Length, WL

The $(WL_{\{e,w\}})$ adds the absolute change between consecutive samples of the signal. This formula does not measure amplitude alone, but rather the amount of accumulated variation within the interval. A high WL may occur when there are rapid oscillations, abrupt contractions, spasms, movement artefacts, or contact noise. Therefore, WL complements RMS and MAV: it helps to distinguish a weak but stable signal from an irregular signal that requires quality review:

(5)

$$W_{L_{e,w}} = \sum_{t=1}^{N_w-1} |s_{e,w}(t+1) - s_{e,w}(t)|$$

2.4.4 Normalization of the Z-score

The value (z_{ij}) is obtained by subtracting the mean (μ_j) of characteristic (j) and dividing it by its standard deviation (σ_j) . This operation centers the data around zero and expresses them in units of standard deviation. Its importance lies in preventing a characteristic with a larger numerical scale from dominating the analysis. After normalization, RMS, MAV, WL, or other variables can be compared within the CART tree or by grouping under more balanced conditions:

(6)

$$Z_{ij} = \frac{X_{ij} - \mu_j}{\sigma_j}$$

Table 4 summarizes the function of each parameter:

Table 4. Variables of sEMG features and normalization

Symbol	Description	Interpretation
RMS_{e,w}	Root mean square of channel e in window w.	Average energy of muscle activation.
MAV_{e,w}	Mean absolute value of channel e in window w.	Unsigned average amplitude.
WL_{e,w}	Waveform length of channel e in window w.	Accumulated variability between consecutive samples.
N_w	Number of samples contained in window w.	Depends on window size and sampling frequency.
x_{ij}	Value of observation i in feature j.	Value before normalization.
mu_j	Mean of feature j.	Distribution center.
sigma_j	Standard deviation of feature j.	Dispersion scale.
z_{ij}	z-score normalised value.	Value expressed in standard-deviation units.

2.5 Decision Tree Model

A decision tree partitions a feature space through sequential rules. In a data-trained tree, each split and threshold is estimated from labelled observations to reduce impurity. ID3 uses information gain, C4.5 uses gain ratio, and CART commonly uses binary splits with Gini impurity and cost-complexity pruning (Breiman et al., 1984; Quinlan, 1986, 1993). The mathematical criteria are included here to formalize how a fitted tree would be optimized. The operational implementation in Section 3.1 is different: it is a rule-based CART-inspired decision system with predefined pilot thresholds, not a CART model trained from the present single-patient archive. This distinction is maintained throughout the revised manuscript.

2.5.1 Entropy H(S)

Entropy measures the uncertainty of the classes within a node. If the observations at the node belong to many classes in similar proportions, H(S) will be high; if almost all belong to a single class, H(S) will be low. In ID3-type trees, the best partition is the one that most reduces this uncertainty. In the sEMG context, a low entropy after dividing by RMS, MAV, or WL indicates that the rule more clearly separates states of weakness:

$$H(S) = - \sum_{c=1}^c p_c \log_2(p_c) \quad (7)$$

2.5.2 Gini Index G(S)

The Gini index estimates the impurity of a node from the class proportions p_c . A pure node has G(S) close to zero, and a data-trained CART model selects the threshold that maximizes impurity reduction. In the present pilot system, the threshold $\text{RMS_03-04}_w \leq 60$ is not learned by minimizing Gini impurity; it is an exploratory rule derived from the descriptive range of the legacy record and is shown only as an example of the binary partition that a CART structure can express:

$$G(S) = 1 - \sum_{c=1}^c p_c^2 \quad (8)$$

2.5.3 Impurity Reduction $\Delta I(S, A)$

This expression compares the impurity of the original node with the weighted impurity of the subsets generated by attribute A. In a fitted tree it quantifies the improvement produced by a candidate split. For the variables used in this study, a future data-trained CART model could compare candidate splits in RMS_03-04_w , D_w or WL_z_w ; the current rule-based implementation does not estimate these thresholds from impurity reduction:

$$\Delta I(S, A) = I(S) - \sum_{v \in \text{Values}(A)} \frac{|S_v|}{|S|} I(S_v) \quad (9)$$

2.5.4 Profit ratio GR(S,A)

The gain ratio adjusts the impurity reduction with the partition information. Its function is to reduce the bias towards attributes with many possible values. Although CART works primarily with Gini and binary partitions, this formula is included because it allows us to contrast the approach with C4.5 and explain why an attribute with many categories is not necessarily better if it fragments the data too much without providing clinically relevant separation:

$$GR(S, A) = \frac{\Delta I(S, A)}{- \sum_{v \in \text{Values}(A)} \frac{|S_v|}{|S|} \log_2 \left(\frac{|S_v|}{|S|} \right)} \quad (10)$$

2.5.5 Binary partitioning by threshold θ

The expression $S_{\leq \theta} = \{x_i \in S : x_{ij} \leq \theta\}$ and $S_{> \theta} = \{x_i \in S : x_{ij} > \theta\}$ defines a binary split of a continuous variable. In a fitted CART model, θ is estimated from training data. In the rule-based pilot used here, θ is predefined for auditability; examples include $RMS_{03-04,w} \leq 60$, $RMS_{03-04,w} \leq 71$, $D_w \geq 0.65$ and $WL_{z,w} > 1.20$. These values are exploratory and are not presented as universal clinical cut-offs:

$$S_{\leq \theta} = \{x_i \in S : x_{ij} \leq \theta\}, \quad S_{> \theta} = \{x_i \in S : x_{ij} > \theta\} \quad (11)$$

2.5.6 Class Prediction $\hat{y}(x)$

The predicted class is assigned by taking the class with the highest probability within the leaf L_m reached by observation x . In practical terms, a sEMG window traverses the tree from the root to a leaf; the leaf summarizes the evidence accumulated by the rules. Thus, the prediction is not opaque: it can be explained by the specific path followed by the window:

$$\hat{y}(x) = \arg \max_{c \in \mathcal{Y}} P(Y = c | x \in \mathcal{L}_m) \quad (12)$$

2.5.7 Empirical Error $R(T)$

The empirical error calculates the proportion of observations misclassified by the T-tree. The indicator function $1(\hat{y}_i \neq y_i)$ is one when the prediction does not match the actual label and zero when it does. This measure allows for the evaluation of the tree, but it should be interpreted with caution in pilot studies, because a low error in a small dataset may reflect overfitting and not necessarily clinical generalizability:

$$R(T) = \frac{1}{|S|} \sum_{i=1}^{|S|} 1(\hat{y}_i \neq y_i) \quad (13)$$

2.5.8 Cost-complexity $R_\alpha(T)$

Cost-complexity pruning adds the tree error and a penalty $\alpha|T|$ multiplied by the number of leaves or tree size. If α is larger, simpler trees are favored; if α is smaller, the tree can grow larger. This formula is key in CART because it avoids generating overly specific rules for a session, sensor, or data capture day. In clinical applications, a simpler tree is usually preferable if it maintains interpretability and stability:

$$R_\alpha(T) = R(T) + \alpha|T| \quad (14)$$

2.5.9 Normalised Weakness Index D_w

The index $D_w = 1 - (RMS_w - RMS_{min}) / (RMS_{max} - RMS_{min})$ rescales the observed RMS activation in the w window using the minimum and maximum values from the reference set. If RMS_w approaches RMS_{max} , the ratio approaches one and D_w approaches zero, suggesting relatively preserved activation. If RMS_w approaches RMS_{min} , the ratio approaches zero and D_w approaches one, suggesting greater relative weakness. Therefore, D_w does not measure clinical diagnosis on its own; it measures a relative position of activation within the observed range:

$$D_w = 1 - \frac{RMS_w - RMS_{min}}{RMS_{max} - RMS_{min}} \quad (15)$$

Table 5 summarizes the function of each parameter:

Table 5. Variables of decision tree formulas

Parameter	Description	Use in the tree
S	Set of observations available at a node.	Basis for measuring impurity.
C	Number of possible severity or muscle-state classes.	Defines the class summation.
p _c	Proportion of class c observations within S.	Input for entropy and Gini.
H(S)	Entropy of set S.	Measures uncertainty; a lower value means a purer node.
A	Candidate attribute for node splitting.	Variable evaluated by the tree.
S _v	Subset of S with value v of attribute A.	Branch produced by the split.
Gain(S,A)	Information gain.	Classical ID3 criterion.
SplitInfo(S,A)	Split information.	Penalizes overly fragmented splits.
GainRatio(S,A)	Gain ratio.	Criterion used by C4.5.
Gini(S)	Gini impurity.	Criterion used by CART.
Delta Gini	Reduction in Gini impurity.	Selects the best binary split.
T	Candidate tree.	Decision structure.
T	Number of terminal nodes.	Tree complexity.
R(T)	Empirical error or classification cost.	Fit component.
alpha	Complexity-penalty parameter.	Controls CART pruning.
D _w	Operational weakness index.	Normalises relative deterioration.
Y	Severity label or operational class.	Target variable for classification.

The D_w value should be analyzed by date, movement phase, and sensor. A single reading may be affected by temporary fatigue, poor device placement, or incomplete fist closure. Conversely, a sustained upward trend in D_w, along with decreasing RMS and MAV under comparable conditions, may serve as an alert for physiotherapy review. For this reason, the index is integrated into the CART tree as a supporting variable and not as the sole classification criterion, see Table 6.

Table 6. Operational interpretation of the D_w index

D _w range	Operational reading	Suggested action
Near 0	Relative activation is conserved within the observed range.	Continue monitoring and compare with baseline.
Intermediate	Possible moderate decrease or transient variation.	Review consistency across sensors and sessions.
Near 1	Lower relative activation or increased operational weakness.	Validate with repeated capture and specialized assessment.

2.6 Target variable and severity rules

A labelled target variable Y is required to train and validate a statistical CART model. In the present rule-based implementation, the ordinal Medical Research Council scale is used only as a conceptual reference for clinically interpretable severity ordering (Medical Research Council, 1943); no claim is made that the pilot thresholds were optimized against MRC labels. Future data-trained validation should use independently assigned clinical or physiotherapeutic labels and report their class distribution, temporal independence and inter-rater procedure.

Table 7. Ordinal reference scale for functional severity

Class	Operational description	Use in the model
Grade 0	No observable muscle contraction or movement.	Maximum weakness class.
Grade 1	Trace contraction without associated movement.	Minimal activation.
Grade 2	Horizontal movement without overcoming gravity.	Severe weakness.
Grade 3	Movement against gravity without external resistance.	Moderate weakness.
Grade 4	Movement against external resistance.	Mild weakness.
Grade 5	Expected normal strength.	Conserved activation.

In the context of the study, the tree can establish rules such as: if the RMS of specific sensors decreases relative to the baseline and the D_w index increases, then a class of greater weakness is assigned; if abrupt amplitudes above the typical

range appear during contraction, a possible spasm or atypical event requiring review is flagged. The workflow consists of the following stages:

- Capture sEMG signals from the eight channels with a sampling frequency of 200 Hz.
- Segment the signals into comparable windows by movement phase and session date.
- Calculate RMS, MAV, WL, and, when relevant, frequency characteristics such as MNF or MDF.
- Normalise features using parameters estimated exclusively from the baseline/reference training subset.
- Apply exploratory clustering only as descriptive analysis of groupings and outliers.
- For future labelled datasets, train and validate ID3, C4.5 or CART models with session- or patient-grouped resampling; the current study applies the predefined rule-based tree.
- Interpret rules, thresholds and decision pathways with the support of clinical specialists or physiotherapists.

3 Results

The legacy visualizations show that many recorded device values lie within an approximately -60 to +60 amplitude band, whereas some contraction segments approach lower absolute amplitudes, see Figure 4. This range is descriptive of the available single-patient record and must not be interpreted as a population 'normal' interval. Longitudinal interpretation therefore relies on within-subject changes in RMS, MAV and D_w under comparable acquisition conditions rather than on a universal raw-amplitude threshold.




Fig. 4. Eight-channel sEMG traces from a sample acquisition. Each panel represents one Myo channel; the horizontal dimension is sample order and the vertical dimension is recorded device amplitude.

Historical comparison is central because ALS progression cannot be inferred from an isolated sample. Figure 5 illustrates the date-range query used to retrieve stored records, and Figure 6 shows channel-wise distributions from the legacy export. Sensors 03 and 04 were prioritized in the pilot rules because they appeared comparatively stable in the recorded task; this channel preference is descriptive and requires confirmation across sessions and additional participants.




Fig. 5. Historical sEMG query interface used to retrieve records by date range for longitudinal comparison; plotted traces correspond to stored sample sequences returned by the selected period.




Fig. 6. Channel-wise distribution of recorded sEMG values in the legacy export. Panels summarize the empirical distribution for individual channels and are descriptive rather than population reference ranges.

In a fitted decision tree, continuous thresholds are estimated from labelled training data. In the present study, the thresholds are predefined exploratory rules selected to make the legacy signal interpretation explicit and auditable. They are therefore described as a rule-based CART-inspired system rather than as a data-trained CART classifier. The operational variables are standardized as RMS_03-04,w (average RMS of channels 03 and 04), D_w (normalized weakness index) and WL_z,w (z-score normalized waveform length).

3.1 Operational rule-based CART-inspired decision system

The operational decision system classifies each valid contraction window into an interpretable monitoring category. The input vector is $X_w = [\text{RMS}_{03-04,w}, D_w, WL_{z,w}]$. Unlike a trained CART estimator, the present tree uses predefined thresholds and deterministic if-then logic. Its purpose is to expose a pilot decision pathway that can later be recalibrated or replaced by thresholds learned from a sufficiently large, clinically labelled dataset. The categories S0-S1, S1, S2, S3 and S4 are operational outputs for review, not diagnoses.

The thresholds are pilot criteria, not learned parameters. $RMS_03-04,w \leq 60$ is anchored to the descriptive concentration of many legacy raw values within approximately ± 60 ; the secondary boundary at 71 preserves a wider exploratory tolerance used in the original operational logic. Because D_w is normalized to $[0,1]$, the cut-points 0.35 and 0.65 define lower, intermediate and upper reference zones around the central range rather than clinically validated grades. $WL_z,w > 1.20$ flags variability more than 1.2 standard deviations above the reference mean and is intentionally used as a sensitivity-oriented review threshold. None of these values is a universal ALS cut-off. They must be recalibrated with additional patients, temporally independent sessions and external clinical labels before performance claims are made, see Table 8.

Table 8. Operational development of the rule-based decision system based on RMS_03-04,w , D_w and WL_z,w

Step	CART operation	Applied criterion	Expected result
1	Validate contraction window	Use only windows associated with the fist-closure phase and with complete date, sensor and repetition metadata.	Avoids classifying resting segments, incomplete captures or records without traceability.
2	Build input vector	$X_w = [RMS_03-04,w, D_w, WL_z,w]$.	Reduces the multichannel signal to three interpretable variables for classification.
3	Evaluate root node	$RMS_03-04,w \leq 60$.	Separates low muscular activation from preserved or elevated activation.
4	Evaluate relative weakness	$D_w \geq 0.65$ or $D_w \geq 0.35$, depending on the branch; predefined exploratory zones.	Distinguishes high weakness, moderate weakness or a discordant condition.
5	Evaluate temporal variability	$WL_z,w > 1.20$; predefined review threshold.	Differentiates sustained weakness from an abrupt event, spasm or possible artefact.
6	Assign terminal leaf	Apply the corresponding binary rule.	Obtains an operational output: S0-S1, S1, S2, S3 or S4.
7	Audit and validate	Compare with previous sessions and independent clinical observations; do not use as a diagnosis.	Reduces the risk of interpreting technical noise as a functional change.

Figure 7 presents the corrected deterministic tree. The low-activation branch begins at $RMS_03-04,w \leq 60$ and is refined by D_w and WL_z,w . The intermediate-amplitude branch $60 < RMS_03-04,w \leq 71$ assigns $D_w \geq 0.65$ to S2 (intermediate condition with historical alert), while lower D_w values remain in the conserved/mild branch unless $WL_z,w > 1.20$, which triggers S4 review. Values $RMS_03-04,w > 71$ are treated as S4 because they lie outside the pilot amplitude tolerance. This rule set is fully synchronized with Table 9 and the Python implementation.




Fig. 7. Corrected operational rule-based tree using RMS_03-04,w , D_w and WL_z,w . The thresholds are predefined exploratory criteria and are not learned from the current dataset.

Figure 7 shows the corrected operational structure. RMS_03-04,w is evaluated first, followed by D_w and, where relevant, WL_z,w . The intermediate condition $60 < RMS_03-04,w \leq 71$ with $D_w \geq 0.65$ now maps consistently to S2 in the figure, Table 9 and code. WL_z,w is retained as a review variable because abrupt variability may represent a spasm, sudden

contraction or technical artefact. All terminal outputs are intended for longitudinal review rather than autonomous clinical classification.

Table 9. Detailed CART decision rules for forming the muscle-weakness

Output	CART rule	Functional interpretation	Suggested action
S4: high weakness with abrupt event	$RMS_{03-04,w} \leq 60$; $D_w \geq 0.65$; $WL_{z,w} > 1.20$	Low activation and high relative weakness coexist with marked variability; this may reflect a spasm, abrupt contraction, fatigue or artefact.	Repeat capture, verify armband contact and compare with clinical/physiotherapy observation.
S3: sustained high weakness	$RMS_{03-04,w} \leq 60$; $D_w \geq 0.65$; $WL_{z,w} \leq 1.20$	Low activation and high relative weakness persist without marked abrupt variability.	Record a longitudinal alert and contrast with clinical strength assessment.
S2: moderate weakness	$RMS_{03-04,w} \leq 60$; $0.35 \leq D_w < 0.65$	Low amplitude with an intermediate normalized weakness index.	Compare with previous sessions and review consistency across repetitions.
S1: discordant low activation / mild weakness	$RMS_{03-04,w} \leq 60$; $D_w < 0.35$	Local RMS is low but the historical weakness index does not support a high-weakness interpretation.	Review armband alignment, channel consistency and repeated capture.
S2: intermediate condition with historical alert	$60 < RMS_{03-04,w} \leq 71$; $D_w \geq 0.65$	Amplitude remains within the wider pilot tolerance, but the historical index suggests relative loss.	Contrast with baseline, session date and independent clinical assessment.
S0-S1: preserved or mild activation	$60 < RMS_{03-04,w} \leq 71$; $D_w < 0.65$; $WL_{z,w} \leq 1.20$	Amplitude remains within the pilot tolerance without a high historical weakness index or marked variability.	Continue longitudinal monitoring under the same protocol.
S4: atypical event / outside pilot tolerance	$RMS_{03-04,w} > 71$ OR $(60 < RMS_{03-04,w} \leq 71$; $D_w < 0.65$; $WL_{z,w} > 1.20)$	Amplitude exceeds the pilot tolerance or marked variability appears in the intermediate-amplitude branch.	Flag for signal-quality and clinical review; do not interpret as isolated improvement or deterioration.

The corrected rule set provides a graded and auditable output while preserving the distinction between sustained low activation and records that require quality review. S3 is the principal sustained-weakness alert; S4 is deliberately conservative because it may represent either a physiologically abrupt event or a technical artefact. S2 includes both moderate weakness in the low-RMS branch and the intermediate-amplitude condition accompanied by a high historical D_w value.

Session-level interpretation should aggregate only valid windows and report the number of candidate, excluded and retained windows for each repetition. Repeated S2/S3 outputs across successive dates may motivate clinical review, whereas multiple S4 outputs should first trigger a capture-quality assessment. The decision system is a transparent monitoring aid and does not replace medical or physiotherapeutic assessment.

3.2 Python implementation of the rule-based decision system

This section provides a reproducible Python implementation of the corrected deterministic rules in Figure 7 and Table 9. Python identifiers `rms_03_04`, `d_w` and `wl_z` correspond exactly to the mathematical variables $RMS_{03-04,w}$, D_w and $WL_{z,w}$. The implementation does not fit a CART estimator from labelled data; it applies the predefined pilot rules to one valid contraction window per row. If an adequately sized and independently labelled dataset becomes available, a separate training pipeline should estimate splits from data and validate them using temporally or patient-grouped resampling.

Code 1. Rule-based CART-inspired decision system in Python

```
# Python implementation of the rule-based CART-inspired decision system
# Required input columns: rms_03_04, d_w, wl_z
# Optional validation column: y_true with labels: S0-S1, S1, S2, S3, S4

import pandas as pd

LABELS = ["S0-S1", "S1", "S2", "S3", "S4"]

def cart_output(rms_03_04: float, d_w: float, wl_z: float) -> str:
    """Classify one contraction window according to Figure 7 and Table 9.

    RMS_03-04,w: average RMS from sensors 03 and 04.
    D_w: normalised weakness index.
    WL_z,w: z-score normalised waveform-length variability.
    """
    if rms_03_04 <= 60:
        if d_w >= 0.65:
            if wl_z > 1.20:
                return "S4" # high weakness with abrupt event
```

```

    return "S3"      # sustained high weakness
    if d_w >= 0.35:
        return "S2"      # moderate weakness
    return "S1"      # discordant low activation or mild weakness

    if rms_03_04 <= 71:
        if d_w >= 0.65:
            return "S2"      # intermediate condition with historical alert
        if wl_z > 1.20:
            return "S4"      # atypical variability in intermediate-amplitude branch
        return "S0-S1"      # preserved activation or mild weakness

    return "S4"      # amplitude outside pilot tolerance

def apply_cart(df: pd.DataFrame) -> pd.DataFrame:
    """Apply CART rules to all windows in a dataframe."""
    required = {"rms_03_04", "d_w", "wl_z"}
    missing = required.difference(df.columns)
    if missing:
        raise ValueError(f"Missing required columns: {missing}")
    df = df.copy()
    df["y_pred"] = df.apply(
        lambda r: cart_output(r["rms_03_04"], r["d_w"], r["wl_z"]), axis=1
    )

    return df

```

The classification routine must be applied only to valid contraction windows. Resting windows, windows with missing or non-finite samples, records without date/repetition metadata, and segments with documented transmission failure are excluded before feature calculation. Quality exclusions must be counted and reported. Normalization parameters are obtained from the baseline/reference training subset only, as defined in Section 2.2.

Code 2. Generic confusion-matrix metrics routine for future labelled validation

```

# Explicit calculation of confusion-matrix metrics for CART outputs
import numpy as np
import pandas as pd
from sklearn.metrics import confusion_matrix, matthews_corcoef

LABELS = ["S0-S1", "S1", "S2", "S3", "S4"]

def calculate_confusion_metrics(y_true, y_pred, labels=LABELS):
    cm = confusion_matrix(y_true, y_pred, labels=labels)
    N = cm.sum()
    rows = []

    for k, label in enumerate(labels):
        TP = cm[k, k]
        FN = cm[k, :].sum() - TP
        FP = cm[:, k].sum() - TP
        TN = N - TP - FP - FN

        precision = TP / (TP + FP) if (TP + FP) else 0.0
        recall = TP / (TP + FN) if (TP + FN) else 0.0
        specificity = TN / (TN + FP) if (TN + FP) else 0.0
        fpr = FP / (FP + TN) if (FP + TN) else 0.0
        fnr = FN / (FN + TP) if (FN + TP) else 0.0
        npv = TN / (TN + FN) if (TN + FN) else 0.0
        f1 = (2 * precision * recall / (precision + recall)
              if (precision + recall) else 0.0)
        support = TP + FN

        rows.append({
            "output": label, "TP": TP, "FP": FP, "FN": FN, "TN": TN,
            "precision": precision, "recall_sensitivity": recall,
            "specificity": specificity, "FPR": fpr, "FNR": fnr,
            "NPV": npv, "F1_score": f1, "support": support
        })

    metrics = pd.DataFrame(rows)
    accuracy = np.trace(cm) / N if N else 0.0
    error_rate = 1.0 - accuracy

```

```

weights = metrics["support"] / N if N else 0.0

summary = {
    "accuracy": accuracy,
    "error_rate": error_rate,
    "macro_precision": metrics["precision"].mean(),
    "macro_recall_balanced_accuracy": metrics["recall_sensitivity"].mean(),
    "macro_specificity": metrics["specificity"].mean(),
    "macro_F1": metrics["F1_score"].mean(),
    "weighted_precision": np.sum(metrics["precision"] * weights),
    "weighted_recall": np.sum(metrics["recall_sensitivity"] * weights),
    "weighted_specificity": np.sum(metrics["specificity"] * weights),
    "weighted_F1": np.sum(metrics["F1_score"] * weights),
    "MCC": matthews_corrocoef(y_true, y_pred)
}
return cm, metrics, summary

```

The current single-patient legacy archive does not provide a documented independent labelled validation dataset with recoverable sample size, class distribution and temporal independence. Consequently, the previously reported summary classification metrics are withdrawn from the revised manuscript because they cannot be rigorously supported without dataset provenance. Code 2 is retained only as a generic routine for calculating metrics when y_true labels and a properly separated validation set are available.

Table 10. Validation status and requirements for future performance reporting

Validation item	Current status	Requirement for a confirmatory study
Independent labelled validation set	Not available in the supplied legacy archive.	Use temporally held-out sessions and, preferably, external patients with independent clinical labels.
Number of validation windows	Not verifiable from the supplied legacy archive.	Report total windows and counts after quality exclusions.
Class distribution	Not available.	Report support for S0-S1, S1, S2, S3 and S4, or the clinically defined target classes.
Temporal/patient independence	Not demonstrated.	Split by session or patient; do not randomly mix overlapping windows from the same recording across train and test sets.
Cross-validation	Not performed for the current rule-based system.	For a data-trained model, use grouped cross-validation with preprocessing fitted within each training fold.
Accuracy / balanced accuracy	Not reported in this revision.	Report only from a documented independent or grouped validation procedure.
MCC and class-wise metrics	Not reported in this revision.	Report MCC, precision, recall, specificity, F1 and support with confidence intervals where feasible.

4 Discussion

The proposed methodology integrates acquisition, structured storage, time-domain feature extraction and an interpretable decision layer for longitudinal sEMG monitoring. The principal contribution is therefore methodological: it defines an auditable chain from a motor task to an operational review category. This is especially important in sEMG because signal amplitude and morphology depend on electrode placement, skin condition, muscle geometry, movement, fatigue and contact quality (De Luca, 1997; Hermens et al., 2000; Merletti & Parker, 2004). The revised manuscript also makes explicit that the current decision tree is rule-based and CART-inspired rather than data-trained. The Gini, entropy and pruning expressions describe the formal decision-tree framework, but the thresholds used in the operational tree were not estimated from a labelled training set. Accordingly, the tree provides transparent pilot rules for within-subject review and not a validated classifier for ALS severity.

Several limitations materially constrain interpretation. First, the study includes one patient and the legacy material available for revision does not contain a complete export from which the exact number of sessions, valid repetitions, retained windows or their full temporal distribution can be reconstructed. Second, the Myo Armband samples EMG at 200 Hz, which is practical for time-domain features but lower than many clinical EMG systems and restricts fine spectral analysis (Phinyomark et al., 2018). Third, the home setting introduces variability in placement, posture, prior fatigue and connection quality. Fourth, the pilot thresholds $RMS_{03-04,w} = 60/71$, $D_w = 0.35/0.65$ and $WL_{z,w} = 1.20$ are exploratory and have not been externally calibrated. Finally, no independent labelled validation set is available; therefore, unsupported performance metrics have been removed.

A confirmatory study should prospectively export window-level data with session identifiers, dates, repetition numbers, quality exclusions and independent clinical labels. Any data-trained CART model should estimate preprocessing parameters and split thresholds only from training data, use session- or patient-grouped cross-validation, and reserve temporally independent or external data for final evaluation. Such validation is required before operational categories are interpreted as clinically meaningful decision support. The present framework should therefore be viewed as a reproducible exploratory protocol and a basis for subsequent multi-patient clinical validation.

5 Conclusions

This study presents a portable, single-patient exploratory framework for longitudinal monitoring of ALS-related muscle activity using eight-channel sEMG and interpretable data-mining rules. The revised methodology standardizes the analytical variables RMS_03-04,w, D_w and WL_z,w, specifies windowing and quality-control requirements, and distinguishes a predefined rule-based CART-inspired system from a data-trained CART model. The corrected Figure 7, Table 9 and Python implementation now use the same terminal outputs for every rule, including the intermediate condition $60 < \text{RMS_03-04,w} \leq 71$ with $D_w \geq 0.65$, which maps to S2. The thresholds remain exploratory pilot criteria and must not be treated as universal ALS cut-offs. Because the source archive does not support a complete reconstruction of session/window counts or an independent labelled validation set, classification-performance claims have been removed. The 200 Hz acquisition rate is suitable for the selected time-domain descriptors but limits fine spectral analysis. Future work must include prospectively documented multi-patient datasets, grouped validation and independent clinical assessment before the framework can support validated clinical decision-making.

References

- Breiman, L., Friedman, J. H., Olshen, R. A., & Stone, C. J. (1984). *Classification and regression trees*. Wadsworth International Group.
- De Luca, C. J. (1997). The use of surface electromyography in biomechanics. *Journal of Applied Biomechanics*, 13(2), 135–163. <https://doi.org/10.1123/jab.13.2.135>
- Han, J., Kamber, M., & Pei, J. (2012). *Data mining: Concepts and techniques* (3rd ed.). Morgan Kaufmann. <https://doi.org/10.1016/C2009-0-61819-5>
- Hermens, H. J., Freriks, B., Disselhorst-Klug, C., & Rau, G. (2000). Development of recommendations for SEMG sensors and sensor placement procedures. *Journal of Electromyography and Kinesiology*, 10(5), 361–374. [https://doi.org/10.1016/S1050-6411\(00\)00027-4](https://doi.org/10.1016/S1050-6411(00)00027-4)
- Kotsiantis, S. B. (2013). Decision trees: A recent overview. *Artificial Intelligence Review*, 39(4), 261–283. <https://doi.org/10.1007/s10462-011-9272-4>
- Medical Research Council. Nerve Injuries Committee. (1943). *Aids to the investigation of peripheral nerve injuries* (2nd rev. ed.; War Memorandum No. 7). H.M. Stationery Office.
- Merletti, R., & Parker, P. A. (Eds.). (2004). *Electromyography: Physiology, engineering, and noninvasive applications*. Wiley-IEEE Press. <https://doi.org/10.1002/0471678384>
- Phinyomark, A., Phukpattaranont, P., & Limsakul, C. (2012). Feature reduction and selection for EMG signal classification. *Expert Systems with Applications*, 39(8), 7420–7431. <https://doi.org/10.1016/j.eswa.2012.01.102>
- Phinyomark, A., Khushaba, R. N., & Scheme, E. (2018). Feature extraction and selection for myoelectric control based on wearable EMG sensors. *Sensors*, 18(5), Article 1615. <https://doi.org/10.3390/s18051615>
- Quinlan, J. R. (1986). Induction of decision trees. *Machine Learning*, 1(1), 81–106. <https://doi.org/10.1007/BF00116251>
- Quinlan, J. R. (1993). *C4.5: Programs for machine learning*. Morgan Kaufmann. <https://doi.org/10.1016/C2009-0-27846-9>
- Witten, I. H., Frank, E., Hall, M. A., & Pal, C. J. (2017). *Data mining: Practical machine learning tools and techniques* (4th ed.). Morgan Kaufmann. <https://doi.org/10.1016/C2015-0-02071-8>
- Yin, R. K. (2018). *Case study research and applications: Design and methods* (6th ed.). SAGE.
- Zhang, Z., Yang, K., Qian, J., & Zhang, L. (2019). Real-time surface EMG pattern recognition for hand gestures based on an artificial neural network. *Sensors*, 19(14), Article 3170. <https://doi.org/10.3390/s19143170>