

Uncertainty-aware PLS regression via neutrosophic representation: Achieving robustness and interpretability in NIR spectroscopy

Ebenezer Aquisman Asare ^a , Dickson Abdul-Wahab ^b, Elsie Effah Kaufmann ^c, Ajay Kanti Das ^d, Rafeah Wahi ^e, Zainab Ngaini ^e, Peter Osei ^f

[Show more](#) 

 Share  Cite

<https://doi.org/10.1016/j.chemolab.2026.105785>

[Get rights and content](#) 

Highlights


- Neutrosophic PLS decomposes observations into Truth, Indeterminacy, Falsity channels.
- Methods achieve 32–40% RMSEP improvement over classical PLS under 30% corruption.
- Clean-data bypass ensures zero penalty when uncertainty channels are negligible.
- Falsity-channel VIP detects corrupted wavelengths with 3.5-fold detection ratio.
- Framework provides traceable quality diagnostics alongside robust calibration.

Abstract

Classical partial least squares (PLS) regression treats all observations and features as equally reliable, making it vulnerable to spike artifacts, detector malfunctions, and heteroscedastic noise commonly encountered in spectroscopic and sensor data. This limitation forces practitioners to choose between applying methods that ignore data quality, risking biased predictions, or manually excluding suspect observations, sacrificing statistical power. This study introduces a neutrosophic framework for PLS regression that explicitly decomposes each observation into Truth, Indeterminacy, and Falsity (T/I/F) channels, enabling uncertainty-aware latent variable modelling. Three variants were developed: NPLS with sample-level reliability weighting, NPLSW incorporating weighted estimation, and PNPLS with element-wise precision weighting through an EM-style procedure. The methods were evaluated using controlled design-of-experiments on IDRC 2016 NIR spectroscopy datasets (MA_A2, A3, B1), with spike corruption ranging from 5% to 30%. Under clean conditions, all methods performed equivalently (RMSEP=0.963), confirming that neutrosophic encoding imposes no penalty when data integrity is maintained. Under severe corruption (30% spikes), neutrosophic variants achieved RMSEP improvements of 32–40% over classical PLS, while maintaining positive R^2 values where PLS collapsed to negative values, indicating complete loss of explanatory power. The falsity-channel VIP exhibited a 3.5-fold detection ratio between corrupted and clean wavelength regions, enabling practitioners to identify problematic spectral regions without distorting genuine predictor importance. This channel-decomposed interpretability provides traceable quality control diagnostics alongside robust predictive calibration, addressing the long-standing robustness-interpretability trade-off in chemometric modelling.

Graphical abstract

Summary of 5-fold cross-validation performance on the IDRC 2016 (B1) NIR spectroscopy dataset under clean conditions and under severe (30%) spike corruption applied to the training data. RMSEP (lower is better) are reported for classical PLS and the PNPLS variant.



[Download: Download high-res image \(549KB\)](#)

[Download: Download full-size image](#)

Introduction

For obvious spike artifacts, experienced practitioners can often identify and mask corrupted wavelengths via visual inspection or simple thresholding. Our contribution is not to replace those established chemometric practices, but to formalize them as an explicit, per-sample/per-wavelength uncertainty representation that can be applied consistently and at scale. In large calibration campaigns, manual screening becomes impractical, and subtle or distributed corruptions (e.g., moderate heteroscedastic noise or small clusters of implausible residuals) can be difficult to spot reliably by eye; the proposed T/I/F encoding provides an automated and auditable way to propagate data quality into both model fitting and interpretation.

A substantial robust PLS literature addresses these failure modes through sample reweighting, robust covariance estimation, or trimming, including partial robust M-regression, least-median and related high-breakdown variants, and boosting-style approaches such as RoBoost-PLSR [1], [2], [3], [4]. These methods can improve stability under contamination, but they often treat data quality through a single set of global or sample-level weights and provide limited diagnostics that localize *where* (which wavelengths) corruption occurred. The neutrosophic formulation used here is complementary: it retains the classical PLS geometry on the Truth channel while exposing wavelength-resolved. Indeterminacy and Falsity signals that can be inspected and used for targeted downweighting without discarding entire spectra.

Spectroscopic, sensor-based, and environmental data frequently suffer from multiple sources of corruption that challenge conventional multivariate regression methods. The data types encountered span near-infrared (NIR) spectra, mass spectrometry records, and sensor measurements from industrial monitoring systems, where the underlying regression problem involves mapping high-dimensional predictor arrays X to response variables y in contexts marked by collinearity and noise. Within these domains, failure modes are well-documented: spike artifacts from detector malfunctions introduce localized high-magnitude deviations, sample-level drift manifests through systematic shifts in baseline or gain, and heteroscedastic noise where measurement precision varies across features or samples remains endemic [1], [2]. Beyond instrumental corruption, missing data mechanisms (missing completely at random, missing at random, or missing not at random) and the presence of genuine outliers further compromise data quality. These failures do not occur uniformly: a single spectrum may exhibit spike artifacts in informative wavelengths while retaining valid signal elsewhere, and noise intensity may be higher in certain spectral regions than others [5]. Classical statistical methods, including ordinary least squares (OLS) and standard partial least squares (PLS), were developed in eras when such heterogeneous corruption was either rare or computationally expensive to model explicitly. Consequently, real-world practitioners face a fundamental tension: apply methods that ignore data quality, risking biased and overconfident predictions, or manually curate and exclude suspect observations, sacrificing statistical power and introducing subjective judgment into analysis pipelines.

Partial least squares regression, formalized through the NIPALS (Nonlinear Iterative Partial Least Squares) algorithm, has proven invaluable for high-dimensional collinear problems by extracting latent variables that maximize covariance between X and y [6], [7]. Nevertheless, standard PLS carries three critical limitations when applied to corrupted or heteroscedastically noisy data. First, the NIPALS algorithm and its variants treat all samples and all features as equally reliable. The standard deflation procedure applies uniform weight to every observation and feature, implicitly assuming that measurement precision and data integrity are homogeneous [8]. When spike artifacts corrupt a minority of wavelengths or a few samples are contaminated with systematic bias, classical PLS incorporates these corrupt points at full weight, allowing their inflated leverage to distort the estimated latent structure [9]. Second, classical PLS offers only weak machinery for handling heteroscedastic noise and mixed data quality. Robust PLS variants address outliers through down-weighting [2], [10], yet existing robust approaches either focus narrowly on sample-level contamination or require separate preprocessing steps that may lose information. Methods such as partial robust M-regression and trimmed or median-based PLS variants have been shown to reduce the influence of extreme points in contaminated calibration problems [1], [3], [4], [9], [11]. In the present study, these robust PLS families were treated as related work and were not implemented as experimental baselines, so empirical comparisons were restricted to classical PLS as the primary reference. Third, classical PLS and standard robust variants provide limited representation of uncertainty itself. A standard PLS model returns point estimates for scores, loadings, and predictions, potentially conveying false confidence when

underlying data quality is poor. Even recent advances in probabilistic and Bayesian PLS variants, while capturing aleatory (data-inherent) uncertainty through noise models, do not systematically distinguish between measured signal, genuine indeterminacy (ambiguous or equivocal information), and detected falsity (clear contamination or implausibility) in a unified framework [12]. This limitation means that many robust formulations still collapsed data quality into a single robust objective or a small set of global weights, which often left limited transparency about where corruption occurred (which wavelengths, which samples) and why a particular observation was effectively ignored. To make uncertainty explicit in a way that supported robustness and diagnosis, this study framed each observation through a neutrosophic representation with three channels: Truth, Indeterminacy, and Falsity. The neutrosophic logic, originating in uncertain reasoning and artificial intelligence, introduces three-valued semantics beyond classical binary truth: a statement or measurement is characterized not only by its truth value (T) but also by an indeterminacy component (I) and a falsity component (F), where $T+I+F$ need not equal unity [13]. Applied to regression, neutrosophic representation assigns each observed datum (or each element within an observation) a three-channel decomposition [12]. The truth channel T captures the deterministic signal or stable component; the indeterminacy channel I quantifies uncertainty, ambiguity, or lack of information; and the falsity channel F represents detected contamination, implausibility, or contradiction with the inferred latent model. In this implementation, the channels were not learned end-to-end; they were constructed deterministically by an encoder from the measured signals. The Truth channel carried the preprocessed spectrum used for learning, while Indeterminacy and Falsity served as bounded, interpretable indicators of uncertainty and contradiction derived from robust statistics of the signal. This separated “what the instrument measured” from “how reliable each region appeared,” enabling uncertainty to be used directly for weighting and interpretation rather than being left implicit in a loss function. On this basis, three PLS-family learners (NPLS, NPLSW, and PNPLS) were evaluated against classical PLS (PLSRegression) as the primary baseline. Robust PLS variants were discussed later to contextualize the contribution, but they were not benchmarked in the present experimental suite. Reliability-weighted NPLS (NPLSW) introduced sample-level weights derived from Indeterminacy and Falsity statistics, so samples flagged as broadly unreliable contributed less during estimation. Probabilistic neutrosophic PLS (PNPLS) moved from sample-level weighting to element-wise precision weighting, treating falsity as a per-cell indicator of heteroscedastic uncertainty and using an EM-style procedure to downweight or effectively impute highly corrupted spectral regions during latent-variable extraction. These design choices were intentionally conservative: the learner operated on the Truth channel used in standard spectroscopy workflows, while uncertainty channels modulated influence rather than redefining the signal itself. This also preserved a clean limiting case, since when Indeterminacy and Falsity were negligible, the weighting mechanisms became uniform and the methods reduced to classical PLS. A controlled design-of-experiments study on the IDRC 2016 NIR datasets (MA_A2), which comprise protein quantification spectra of food and agricultural samples. Systematically varied spike corruption parameters to quantify sensitivity of each method under increasing artifact severity. Two additional IDRC 2016 NIR spectroscopy datasets (A3 and B1) were then used for validation under clean conditions and under a severe spike-corruption regime applied to training data, with performance assessed using RMSEP and R^2 stability. To connect robustness to interpretability, the study also examined VIP behaviour and feature-importance stability, testing whether methods that resisted corruption also maintained coherent and reproducible spectral importance profiles.

This work advanced neutrosophic latent variable modelling through four concrete contributions. First, a neutrosophic observation model for supervised learning was formulated with explicit Truth, Indeterminacy, and Falsity channels, and the learning procedure was kept compatible with high-dimensional, collinear regression by retaining a NIPALS-style alternating optimization structure. Second, two weighted extensions were introduced: reliability-weighted NPLS, which applied sample-level weights derived from estimated or known observation reliability, and precision-weighted PNPLS, which applied element-wise precision weights through a heteroscedastic variance model. Both extensions were designed to integrate naturally into NIPALS deflation while preserving the usual latent-variable interpretation. Third, the study demonstrated substantive robustness gains under controlled synthetic corruption regimes, where spike artifacts were injected at contamination levels spanning 5–30%, and under real NIR spectroscopy datasets, showing that neutrosophic variants maintained lower prediction error and more stable feature importance than classical PLS as data quality degraded. Fourth, VIP-based interpretability tools were provided that decomposed feature importance by channel (T/I/F), allowing practitioners to separate predictive signal from data-quality artifacts and to use the falsity channel as a practical quality-control diagnostic for instrument monitoring and dataset curation. The remainder of the paper was organized as follows. The Methods section defined the neutrosophic observation model, the construction of the T/I/F channels, and the NIPALS-based learning procedures for NPLS, NPLSW, and PNPLS, including the associated prediction rules and the limiting reduction to classical PLS when uncertainty channels were negligible. The Results section first reported controlled spike-corruption experiments (5–30% contamination) and then presented validation on the IDRC 2016 protein datasets (MA_A2, A3, and B1), with performance summarized using RMSEP, R^2 , and robustness trends under degrading signal conditions. The Discussion section interpreted these findings in terms of robustness mechanisms and practical utility, with emphasis on VIP stability and the channel-decomposed interpretability outputs that supported quality-control reasoning alongside predictive calibration.

Section snippets

Notation and problem statement

Let (n) denote the number of samples, (p) the number of predictor variables, and (q) the number of response variables. The observed predictor matrix is ($\mathbf{X} \in \mathbb{R}^{n \times p}$) and the response is either univariate ($\mathbf{y} \in \mathbb{R}^n$) or multivariate ($\mathbf{Y} \in \mathbb{R}^{n \times q}$). The neutrosophic

representation augments each scalar observation with a triplet of channels ($\mathbf{T}, \mathbf{I}, \mathbf{F}$) and is stored as tensors $\mathbf{X}^{(TIF)} \in \mathbb{R}^{n \times p \times 3}$ and $\mathbf{Y}^{(TIF)} \in \mathbb{R}^{n \times q \times 3}$, where the last axis indexes Truth, Indeterminacy, and Falsity. In the implemented encoding, these tensors are ...

NPLS: reliability-weighted neutrosophic NIPALS-PLS

Let ($\mathbf{X}^{(TIF)} \in \mathbb{R}^{n \times p \times 3}$) and ($\mathbf{Y}^{(TIF)} \in \mathbb{R}^{n \times q \times 3}$). In the implemented NPLS-family, learning is performed on an effective “observed” matrix extracted from the neutrosophic tensor. For the experiments executed in `generate_paper_figures.py`, the fusion was deliberately set to Truth-only via **channel weights** = $(1, 0, 0)$, so comparisons isolate reliability mechanisms rather than feature attenuation. The model code similarly operates on the Truth channel in both NPLS and its weighted variant. A general fusion ...

Controlled robustness (DOE on MA_A2)

A design-of-experiments was conducted to systematically assess predictive stability under controlled corruption of real near-infrared (NIR) spectroscopy data (MA_A2). Two corruption parameters were varied across the training set: spike/outlier proportion (p_{out} ranging from 0 to 30 percent) and additive noise level (σ_x from 0.00 to 0.10). Four methods were compared: classical partial least squares (PLS) as the baseline, neutrosophic PLS (NPLS, using reliability-weighted truth values), NPLSW ...

Robustness-interpretability trade-off

Practitioners working with spectroscopic data face a persistent dilemma: robust methods that protect against data corruption often sacrifice transparency, while interpretable methods remain vulnerable to outliers and artifacts. Classical robust approaches such as M-regression, trimmed PLS, and RoBoost-PLSR improve stability under contamination, yet they provide limited insight into why particular observations receive reduced weight [2], [3]. Neutrosophic methods close this gap by making ...

Conclusion

This work demonstrated that neutrosophic representation offers a principled approach to robust PLS regression by making data quality uncertainty explicit through structured Truth, Indeterminacy, and Falsity channels. The proposed methods (NPLS, NPLSW, and PNPLS) achieved substantial robustness gains under spike contamination while maintaining exact equivalence to classical PLS on clean data, as verified both analytically and numerically. The empirical results across three NIR spectroscopy ...

Ethical approval and consent to participate

This paper was designed and written in the National Nuclear Research Institute, Ghana Atomic Energy Commission, Legon-Accra All the images in this research were originally designed by the authors. ...

Consent for publication

Not applicable ...

Funding

This research received no funding ...

CRedit authorship contribution statement

Ebenezer Aquisman Asare: Conceptualization, Data curation, Formal analysis, Investigation, Methodology, Supervision, Validation, Visualization, Writing – original draft, Writing – review & editing. **Dickson Abdul-Wahab:** Data curation, Formal analysis, Software, Writing – original draft. **Elsie Effah Kaufmann:** Resources, Supervision, Visualization, Writing – review & editing. **Ajoy Kanti Das:** Methodology, Visualization, Writing – review & editing. **Rafeah Wahi:** Supervision, Validation, Writing – ...

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper. ...

Acknowledgment

The authors acknowledge the contribution of colleagues from the Nuclear Chemistry and Environmental Research Centre, National Nuclear Research Institute (NNRI), Ghana Atomic Energy Commission (GAEC), Box LG 80, Legon-Accra, Ghana. ...

References (14)

F.B. Lavoie *et al.*

[A novel robust NL-PLS regression methodology](#)

Chemom. Intell. Lab. Syst. (Jan. 2019)

M. Metz *et al.*

[A novel robust PLS regression method inspired from boosting principles: roboost-plsr](#)

Anal. Chim. Acta (Sep. 2021)

S. Serneels *et al.*

[Partial robust M-regression](#)

Chemom. Intell. Lab. Syst. (Oct. 2005)

D. Tanzilli

[Locally-weighted-RoBoost-PLS: a multivariate calibration approach to simultaneously cope with non-linearities and outliers](#)

Anal. Chim. Acta (Aug. 2025)

P. Geladi *et al.*

[Partial least-squares regression: a tutorial](#)

Anal. Chim. Acta (1986)

B. Liebmam *et al.*

[Robust and classical PLS regression compared](#)

J. Chemom. (Mar. 2010)

S. Wold *et al.*

[The collinearity problem in Linear regression. The partial least squares \(PLS\) approach to generalized inverses](#)

SIAM J. Sci. Stat. Comput. (Sep. 1984)

There are more references available in the full text version of this article.

Cited by (0)

[View full text](#)

© 2026 Elsevier B.V. All rights are reserved, including those for text and data mining, AI training, and similar technologies.



All content on this site: Copyright © 2026 Elsevier B.V., its licensors, and contributors. All rights are reserved, including those for text and data mining, AI training, and similar technologies. For all open access content, the relevant licensing terms apply.

