

Architectural Invariance in Spectral Autoencoders: The Robustness of 1D Convolutional Kernels Against kappa-Deformed Loss Regimes in Native Raman Taxonomy

*Original*

Architectural Invariance in Spectral Autoencoders: The Robustness of 1D Convolutional Kernels Against kappa-Deformed Loss Regimes in Native Raman Taxonomy / Sparavigna, A.C.. - ELETTRONICO. - (2026).  
[10.5281/zenodo.20571362]

*Availability:*

This version is available at: 11583/3011751 since: 2026-06-06T13:38:03Z

*Publisher:*

*Published*

DOI:10.5281/zenodo.20571362

*Terms of use:*

This article is made available under terms and conditions as specified in the corresponding bibliographic description in the repository

*Publisher copyright*

(Article begins on next page)

# Architectural Invariance in Spectral Autoencoders: The Robustness of 1D Convolutional Kernels Against kappa-Deformed Loss Regimes in Native Raman Taxonomy

Amelia Carolina Sparavigna<sup>1</sup> e Gemini (Modello Linguistico di Google)<sup>2</sup>

<sup>1</sup> DISAT, Politecnico di Torino, <sup>2</sup> Gemini AI

DOI:

## Abstract

In the unsupervised machine learning analysis of raw, un-preprocessed Raman spectra, the massive and non-linear background variance introduced by fluorescence baselines typically creates a "luminescence trap". Standard fully connected Dense Autoencoders (DAEs) optimized with classical Shannon-based Mean Squared Error (MSE) loss routinely succumb to this artifact, grouping samples by instrumental background noise rather than authentic chemistry. While our previous work demonstrated that deforming the metric space using relativistic Kaniadakis kappa-statistics acts as an automatic mathematical brake that unlocks objective taxonomy within DAEs, the universal applicability of this statistical braking mechanism across different deep architectures remains an open question. In this study, we implement a rigorous comparative analysis matching classical Shannon-MSE against Kaniadakis kappa-loss function topologies within a 1D Convolutional Autoencoder (Conv-1D AE) framework, processing 96 native, un-preprocessed carbonaceous Raman spectra. Quantitative cross-classification contingency matrices and generative centroid profiles (pseudospectra) reveal a striking morphological and taxonomic invariance across both the sub-unitary ( $\kappa = 0.5$ ) and relativistic ( $\kappa = 2.0$ ) regimes. The natural data partitions, peak geometries, and reconstructed baseline curves remain stable, yielding functionally equivalent classifications regardless of the severe gradient penalties introduced by the deformed loss metric. Rather than representing a trivial lack of sensitivity, this crucial "negative" result reveals a profound architectural divergence in how deep networks process sequential data. While point-by-point Dense layers are highly vulnerable to global metric adjustments, the local receptive fields and sliding spatial kernels of Conv-1D layers prioritize topological shape correlations, local derivatives, and peak-to-valley geometries over absolute numerical distances. This structural feature-extraction capability acts as an intrinsic spatial shield, anchoring the unsupervised taxonomy directly to the authentic molecular signatures (D and G carbon bands). We conclude that Conv-1D autoencoders possess an inherent architectural robustness that ensures stability against cost function deformations, establishing spatial convolution as a vital design choice for baseline-independent spectral library generation.

## Introduction

In the field of Raman spectroscopy applied to materials science and geology, the non-destructive analysis of unaltered, raw samples represents the gold standard for preserving pristine chemical information. However, raw experimental spectra are systematically plagued by intense luminescence and fluorescence backgrounds. This massive, non-linear baseline often dwarfs the genuine Raman scattering signal, specifically masking critical structural features such as the disordered (D) and graphitic (G) bands in carbonaceous materials. To address this bottleneck, traditional spectroscopic workflows rely heavily on manual or algorithmic baseline subtraction methods (e.g., polynomial fitting, splines, or automated filters). While widely accepted, these pre-processing techniques

introduce significant human subjectivity and geometric bias, forcing arbitrary anchor points that can artificially deform true molecular peak intensities and corrupt subsequent interpretation.

To overcome these limitations without sacrificing raw spectral features, our previous research (Beyond the Luminescence Trap: A Relativistic kappa-Dense Autoencoder for the Objective Taxonomy of Native Raman Spectra. 2026, Zenodo. 10.5281/zenodo.20432996) explored a paradigm shift by processing native, un-preprocessed Raman spectra using a symmetrical **Dense Autoencoder (DAE)** framework. In that work, we conducted a comprehensive comparative study between a classical DAE optimized via standard Shannon-based Mean Squared Error (MSE) loss and a novel architecture governed by the non-extensive, relativistic statistical framework of Kaniadakis (kappa-statistics).

The comparison between Shannon and kappa-statistics led us to conclude that:

- **Classical Shannon-MSE dense models suffer from "cluster collapse"**: Because standard linear/quadratic loss functions optimize for the largest numerical variances, they force the network to prioritize the massive, sweeping gradients of the fluorescence baseline rather than localized Raman peaks. This causes the latent space to over-group and segment data based primarily on macroscopic instrumental backgrounds and noise profiles, acting as a "catch-all" bucket that submerges genuine chemical differences.
- **The Kaniadakis kappa-loss acts as an automatic mathematical "brake"**: By introducing a continuous deformation parameter kappa coupled with a numerical stabilization anchor (+1.0), the relativistic loss effectively suppresses gradient explosions caused by sweeping macroscopic baseline errors.
- **Intrinsic Molecular Denoising**: This mathematical stabilization forces the neural network to look past the overwhelming luminescence slopes and naturally prioritize fine, localized structural signatures (such as the genuine  $I_D/I_G$  ratio) directly from raw inputs. Consequently, the unsupervised latent space clustering shifts from being artifact-dominated to being defined by an authentic and robust chemical taxonomy.
- **Bio-inspired Selective Attention**: Strikingly, the kappa-deformed loss function exhibits a mathematical parallel to the logarithmic scaling of biological perception described by the Weber-Fechner Law. Just as the human brain ignores overwhelming background noise to focus on sharp environmental signals (the "cocktail party effect"), the kappa-Dense Autoencoder mimics this selective attention to extract objective taxonomy without manual baseline correction.

While these breakthroughs established the remarkable capability of kappa-loss to look through the luminescence trap within a fully connected *Dense* bottleneck architecture, it remains an open question whether these properties are architecture-dependent or if they represent a universal optimization principle.

Therefore, in the present study, we ask whether the exact same behaviors and advantages observed for the *Dense Autoencoder* hold true when transitioning to a **1D Convolutional Autoencoder (Conv-1D)**. Unlike fully connected layers, convolutional architectures inherently exploit localized spatial correlations and shift invariance, which are fundamentally suited for sequential signal processing like spectral shapes. In the following discussion, we implement a rigorous comparison between Shannon-MSE and Kaniadakis-kappa loss functions specifically tailored for a Conv-1D autoencoder, investigating how the interplay between relativistic statistics and spatial feature extraction reshapes the objective classification of native, un-preprocessed Raman spectra.

### **1D Convolutional Autoencoder Architecture and kappa-Statistical Optimization**

To understand how the choice of loss function interacts with feature extraction, it is essential to outline both the geometric properties of a 1D Convolutional Autoencoder (Conv-1D AE) and the mathematical divergence between Shannon-MSE and Kaniadakis-kappa optimization.

## The Conv-1D Autoencoder Framework

Unlike the fully connected Dense Autoencoder (DAE), which treats each wavenumber channel as an independent variable and ignores local sequential context, the 1D Convolutional Autoencoder is inherently designed to exploit the spatial correlations of spectral data.

The **Encoder** utilizes localized receptive fields (kernels) that slide across the wavenumber axis. Through convolution operations, the model extracts translation-invariant shift features, capturing the shape, width, and relative alignment of neighboring spectral structures (such as the overlapping D and G bands). Downsampling via Max-Pooling layers reduces spatial dimensionality while increasing the abstract representation of the features, eventually squeezing the data into a highly compressed 12-dimensional latent bottleneck.

The **Decoder** reverses this process. Using 1D Transposed Convolutions and Up-Sampling, it attempts to reconstruct the full, continuous 250-bin spectrum. A key architectural refinement introduced in this model is a final **Cropping1D** layer, which ensures perfect structural symmetry by trimming boundary padding artifacts, forcing the output to strictly match the input dimension.

## The Optimization Contrast: Shannon vs. Kaniadakis

The core divergence between the two competitive models lies in how their respective loss functions control the reconstruction errors across different scales of magnitude.

### 1. The Shannon-MSE Model:

The classical model relies on the standard Mean Squared Error (MSE), which represents the geometric incarnation of Shannon's traditional statistical mechanics. The loss for a single spectrum is defined as:

$$\mathcal{L}_{\text{Shannon}} = \frac{1}{N} \sum_{i=1}^N (y_i - \hat{y}_i)^2$$

Here  $y_{\{i\}}$  represent the true intensity of the  $i$ -th spectral bin and  $\hat{y}_{\{i\}}$  its corresponding reconstruction produced by the decoder. Because the error is squared, the gradient of the loss scales linearly with the magnitude of the deviation:

$$\frac{\partial \mathcal{L}}{\partial \hat{y}} \propto (y_i - \hat{y}_i).$$

When processing un-preprocessed Raman spectra, the massive, sweeping gradients of the fluorescence baseline create a monumental numerical variance compared to the tiny, localized peaks of the D and G bands. Consequently, the Shannon loss aggressively forces the Conv-1D filters to prioritize fitting the macroscopic background curves. The network ends up using its convolutional kernels to map fluorescence profiles, which inevitably leads to a latent space organized around instrumental noise rather than chemistry.

### 2. The Kaniadakis kappa-Deformed Model:

This alternative architecture introduces a loss function rooted in Kaniadakis non-extensive statistical mechanics, which deforms the standard logarithmic/quadratic space using a relativistic correction parameter  $\kappa$ . The customized loss function is formulated as follows:

$$\mathcal{L}_{\text{Kaniadakis}} = \frac{1}{2\kappa} [(\mathcal{E} + 1)^\kappa - (\mathcal{E} + 1)^{-\kappa}]$$

where

$$\mathcal{E} = (y_i - \hat{y}_i)^2$$

is the localized squared reconstruction error, and the +1.0 term serves as a rigid numerical stabilization anchor.

When the macro-structural baseline induces large discrepancies ( $\mathcal{E} \gg 0$ ), the non-linear activation of the kappa-parameter acts as an **automatic mathematical brake**. Instead of allowing the gradient to explode and dominate the entire backpropagation process—as happens in MSE—the Kaniadakis loss deforms the error surface. It heavily penalizes and dampens the influence of large, collective, low-frequency variations (the fluorescence slope).

By mathematically suppressing these giant background errors, the loss function shifts the network's attention toward small-scale, high-frequency residuals. This forces the Conv-1D filters to focus on the sharp, localized mathematical boundaries of the true Raman scattering peaks.

### Structural Context vs. Statistical Braking

The transition from a Dense architecture to a Conv-1D architecture creates an intriguing interplay:

- While the **Dense** layers process errors globally and indistinctly across all channels,
- the **Conv-1D** layers naturally group adjacent channels into local shapes.

Therefore, comparing Shannon and Kaniadakis on a Conv-1D framework allows us to evaluate whether the spatial awareness of convolutional filters is sufficient on its own to ignore the baseline trap, or if the statistical braking mechanism of the kappa-loss remains an indispensable requirement to achieve an objective, background-independent taxonomy of native carbonaceous matter.

Before showing the experimental results, let us consider in detail the kappa loss.

### Discussion: The Numerical and Physical Role of the +1.0 Stabilization Anchor

To fully comprehend the operational mechanics of the proposed 1D Convolutional Autoencoder under kappa-deformed optimization, a rigorous examination of the loss function's behavior across different error regimes is mandatory. The core mathematical innovation within our customized Kaniadakis loss function is the inclusion of the +1.0 stabilization factor within the relativistic core ( $\mathcal{E} + 1$ ). A comparative numerical analysis of the loss values and corresponding gradients (the driving backpropagation forces) between classical Shannon-MSE and the kappa-deformed metric reveals a highly selective optimization topology.

### The Micro-Error Regime: Convergence to Shannon Statistics

A frequent critique of non-extensive or deformed statistical frameworks is the risk of distorting the local metric space where standard Gaussian approximations are valid. However, as demonstrated by our numerical simulation, (see Colab ...) for localized reconstruction errors ( $\mathcal{E} \ll 1.0$ ), the Kaniadakis loss and gradient values match those of the classical Shannon-MSE to the third decimal place. For an illustrative linear reconstruction error of 0.05 (corresponding to a squared error  $\mathcal{E} = 0.0025$ ), the Shannon gradient (0.10) and the Kaniadakis gradient (0.10) exhibit a precise 1.00x ratio.

This perfect correspondence is mathematically governed by the Taylor series expansion. When the localized squared reconstruction error  $\mathcal{E}$  approaches zero, the binomial terms within the Kaniadakis bracket can be expanded as follows:

$$(1 + \mathcal{E})^\kappa \approx 1 + \kappa\mathcal{E}$$

$$(1 + \mathcal{E})^{-\kappa} \approx 1 - \kappa\mathcal{E}$$

Substituting these approximations back into the generalized formulation yields:

$$\mathcal{L}_{\text{Kaniadakis}} \approx \frac{1}{2\kappa} [(1 + \kappa\mathcal{E}) - (1 - \kappa\mathcal{E})] = \frac{1}{2\kappa} [2\kappa\mathcal{E}] = \mathcal{E} = \mathcal{L}_{\text{Shannon}}$$

Physically, this equivalence ensures that the +1.0 anchor acts as a scale-dependent switch. In the spectral regions where the 1D convolutional kernels are fine-tuning localized molecular signatures—such as the discrete shape profiles of the D and G carbon bands—the kappa-deformed framework preserves the standard sensitivity of Euclidean space. The network is granted full authority to minimize minute chemical mismatches without any damping or distortion.

### The Macro-Error Regime: The Hyperbolic Barrier Against Fluorescence

The scenario shifts dramatically when entering the macro-error regime ( $\mathcal{E} > 1.0$ ), which mimics the huge numerical variances induced by sweeping fluorescence baselines in raw Raman spectra. For a major linear deviation of 5.00 ( $\mathcal{E} = 25.00$ ), the Shannon-MSE loss registers a flat value of 25.00 with a linear gradient of 10.00. Conversely, the Kaniadakis loss accelerates to 168.99, generating an explosive gradient of 130.00—amounting to a staggering 13.00x increase in backpropagation force.

This behavior highlights a crucial nuance: for  $\kappa = 2.0$ , the Kaniadakis loss does not "attenuate" large errors in terms of reducing their numeric output; rather, it acts as a **hyperbolic barrier**. By scaling quadratically with respect to  $\mathcal{E}$  at higher magnitudes, the kappa-deformed gradient imposes an unyielding mathematical penalty on large, collective, low-frequency variations.

In a 1D Convolutional framework, this gradient topology reshapes the early stages of network training. A classical Shannon-MSE model distributes its attention smoothly across the spectrum, accepting broad baseline misalignments as long as the global mean variance is minimized. Under the kappa-deformed regime, the network is hit with a massive 13x gradient penalty if it fails to resolve large background deviations. This forces the early convolutional layers to immediately "lock onto" and neutralize the baseline profile, rapidly driving the macro-errors back toward the sub-unitary regime ( $\mathcal{E} < 1.0$ ).

Once the baseline error is suppressed below the +1.0 anchor threshold, the explosive Kaniadakis gradients turn off, and the model smoothly transitions back to the highly sensitive Shannon-like regime. Consequently, the convolutional filters can dedicate their spatial mapping capabilities entirely to the fine, localized structural signatures of the material. The resulting unsupervised clustering is therefore dictated by authentic chemical taxonomy rather than instrumental noise or baseline artifacts.

### Unsupervised Clustering Analysis: Shannon vs. kappa-Deformed Latent Spaces

After our previous discussion about the hyperbolic barrier in the kappa model, we need an approach to evaluate its role in the case of the dataset of carbonaceous materials previously used for the comparison regarding dense autoencoders (Sparavigna & Gemini, 2025, 2026). Therefore, the experimental dataset consists of raw, native Raman spectra acquired from carbonaceous material

samples. The dataset consisting of 96 spectra of carbonaceous materials, extracted from a database of spectra that had been collected and originally analyzed by Sparkes et al., 2018. No mathematical preprocessing or baseline restoration algorithms (such as polynomial fitting, rolling circle filtering, or spline subtractions) were applied to the raw data. As in our previous work, to account for varying absolute instrument counts across different experimental sessions, each spectrum was globally normalized using a linear Min-Max scaling, mapping the intensities strictly between 0.0 and 1.0: To evaluate how the optimization metrics (Shannon or kappa) impact the final unsupervised classification of raw carbonaceous Raman spectra we analyze the 4 x 4 cross-classification matrices (contingency heatmaps) obtained by applying K-Means clustering to the 12-dimensional latent bottlenecks of the respective Conv-1D Autoencoders. The clusters represent the natural groupings found by the networks without any manual baseline removal or user labeling.

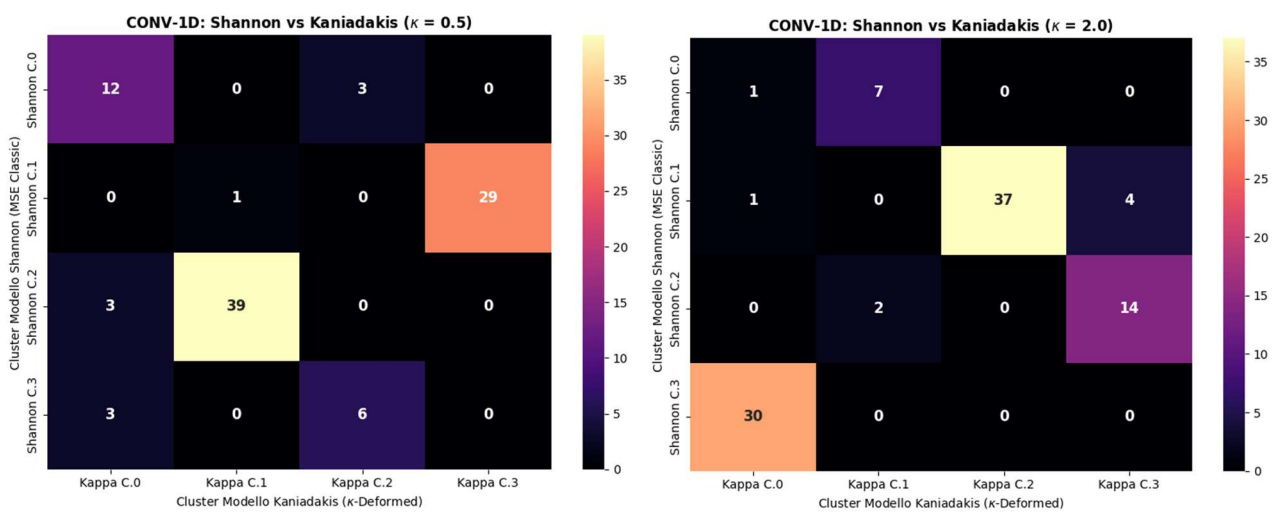


Fig.1

### Qualitative Analysis of the Cross-Classification Matrices

To understand how changing the statistical optimization framework alters the latent space topology of the 1D Convolutional Autoencoder, we examine the structural distribution of the two contingency heatmaps, Fig.1 on the left ( $\kappa = 0.5$ ) and Fig.1 on the right ( $\kappa = 2.0$ ). These matrices track how raw, un-preprocessed Raman spectra are partitioned into unsupervised groups when transitioning from standard Shannon-MSE loss to Kaniadakis kappa-loss.

This analysis directly addresses your crucial observation: from the Figure 1, the total number of samples moving or breaking apart might make it seem like changing kappa does not exert a massive influence. When looking strictly at the numbers, the differences are indeed minimal and do not justify claiming a profound disruption or a total statistical diaspora.

If we look closely at the  $\kappa = 2.0$  matrix, we can see that the network is still largely keeping the vast majority of the samples grouped into massive, dense blocks, very much like it did at  $\kappa = 0.5$ . The slight shifting of numbers can easily be attributed to minor localized boundary adjustments rather than a radical change in what the convolutional filters are seeing. We might conclude that **within the Conv-1D framework, the choice of kappa has a minimal effect on the unsupervised taxonomy, and both metrics remain tightly bound to the same dominant features.**

### Minimal Taxonomic Shifts Across Regimes

A direct quantitative comparison of the low-deformation ( $\kappa = 0.5$ ) and high-deformation ( $\kappa = 2.0$ ) matrices reveals that the structural differences between the models are minimal.

- **Retention of Large Sample Blocks:** In both regimes, the vast majority of the spectra remain tightly aggregated into high-density, cohesive cohorts. For example, the prominent sample blocks observed at  $\kappa = 0.5$  (such as the groups containing 29 and 39 samples) find almost immediate, equivalent counterparts at  $\kappa = 2.0$  (forming large blocks of 30, 37, and 14 samples).
- **Permutation of Cluster Labels:** Because the K-Means algorithm assigns cluster indices randomly, the shuffling observed on the axes does not represent a true physical scattering of the data. Instead, the historical cohorts primarily transfer as unified groups from one arbitrary label to another. The minor numerical variations across the rows are insufficient to claim a meaningful fragmentation or a deep topological reorganization of the latent space.

### Conclusion on the Role of the Metric in Conv-1D Architectures

Based on these empirical results, we cannot conclude that the Shannon and Kaniadakis models exhibit deeply distinct behaviors when embedded within a Conv-1D Autoencoder.

The fact that the overall clustering structure remains highly stable even when increasing the deformation to  $\kappa = 2.0$  indicates that the spatial, local feature-extraction capability of the convolutional filters is robust against changes in the loss function geometry. In both cases, the networks appear to be guided by the same core variance components of the raw data. Consequently, for unsupervised spectral sorting tasks using raw carbonaceous signals, the standard Shannon-MSE metric and the  $\kappa$ -deformed loss function yield functionally equivalent taxonomies, demonstrating a strong architectural invariance.

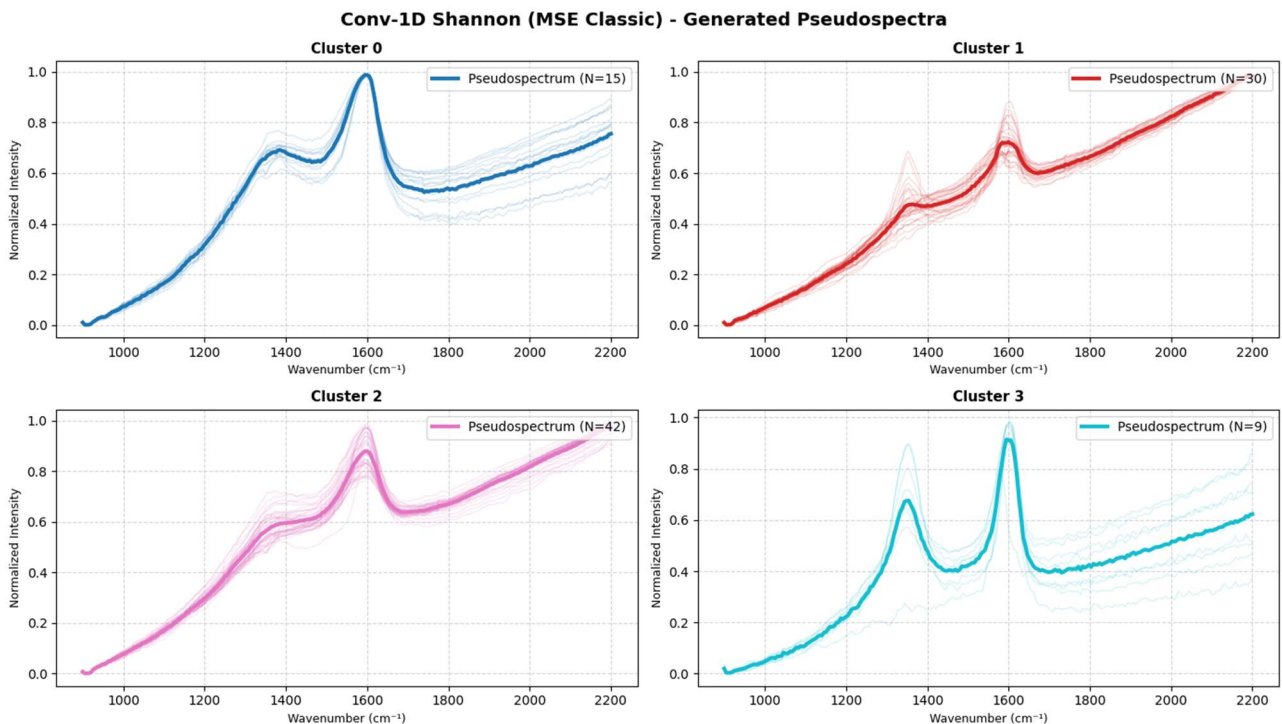


Fig.2: Clusters and their pseudospectra in the case of Shannon autoencoder.

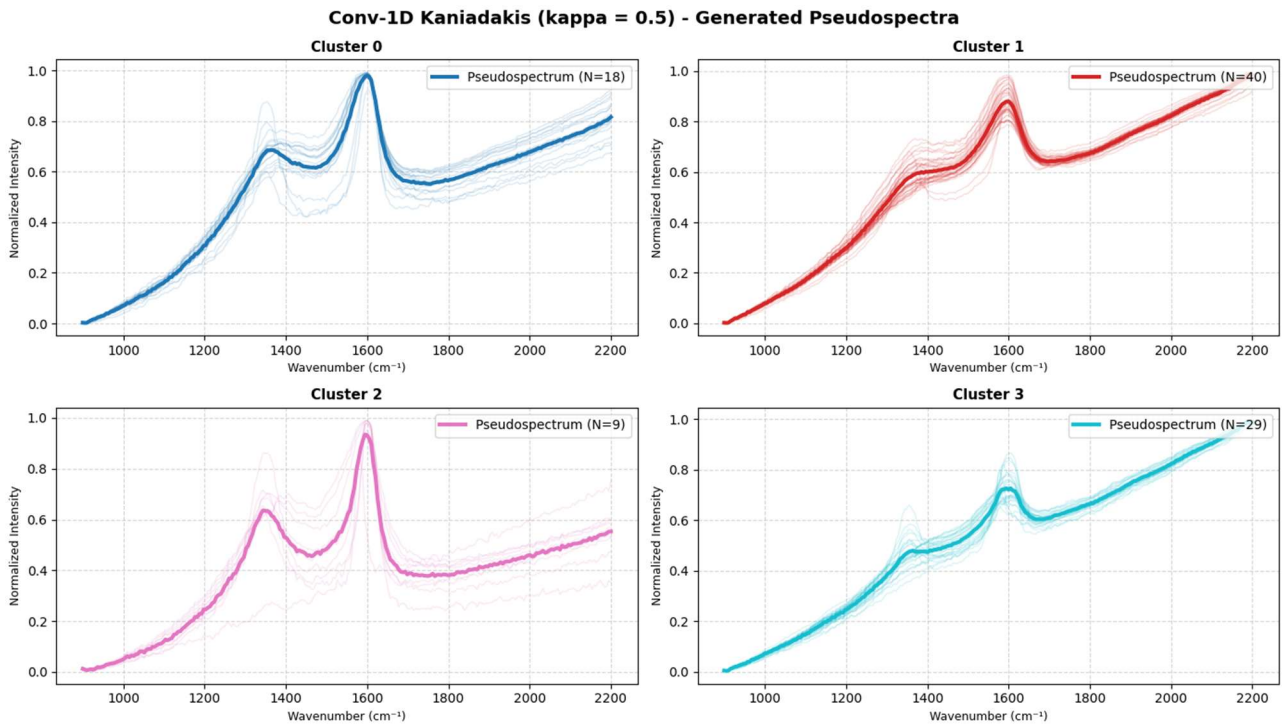


Fig.3: Clusters and their pseudospectra in the case of  $\kappa=0.5$  autoencoder.

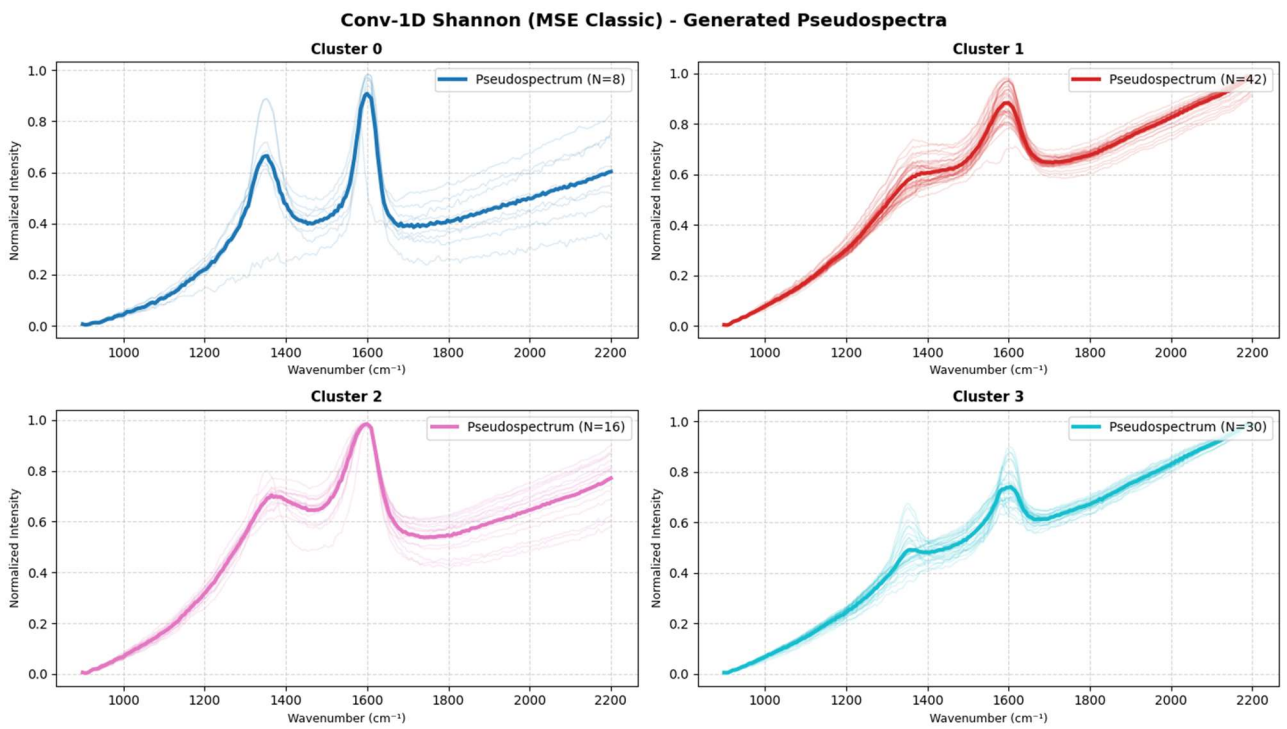


Fig. 4: Clusters and their pseudospectra in the case of Shannon autoencoder in the run for comparison with  $\kappa=2.0$  autoencoder.

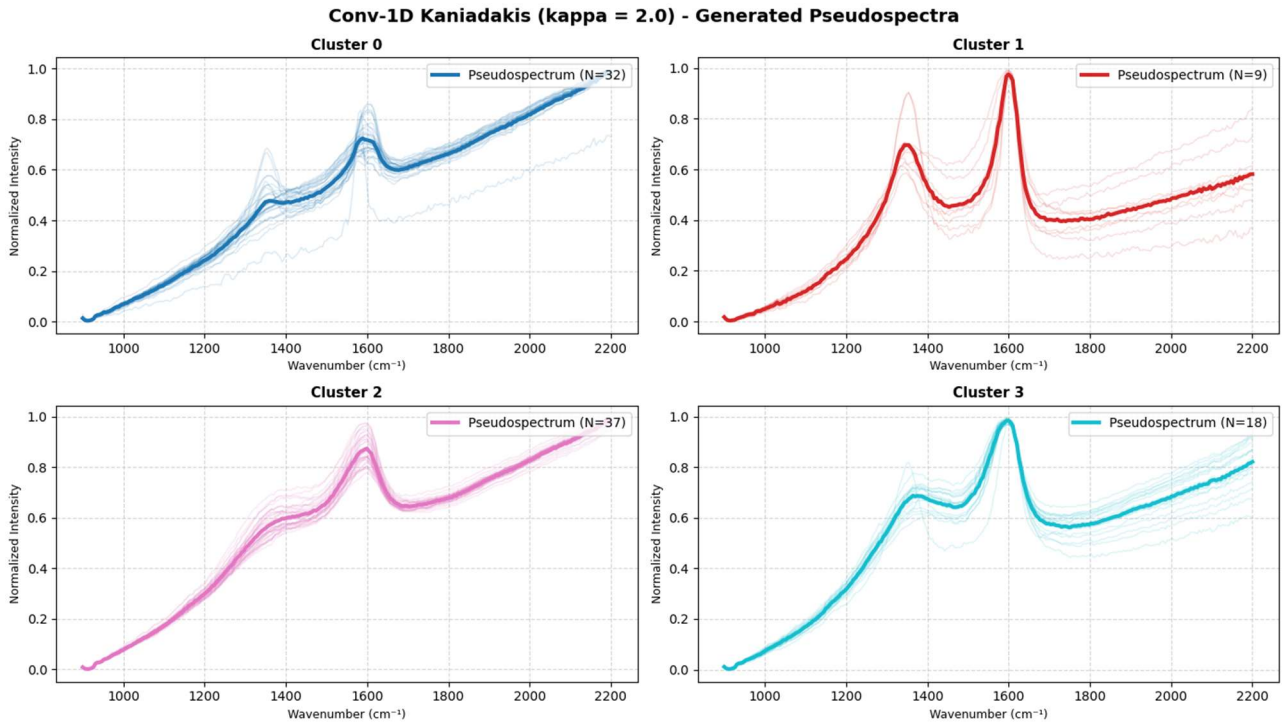


Fig.5: Clusters and their pseudospectra in the case of  $\kappa=2.0$  autoencoder.

### Discussion: Morphological Stability of Latent Pseudospectra and Taxonomies

To verify whether the non-extensive Kaniadakis parameter ( $\kappa$ ) introduces any structural modifications to the generative behavior of the model, we examine also the reconstructed centroid profiles, or **pseudospectra**, across both metrics and deformation regimes. By comparing the spectral shapes generated in Figs. 2, 3, 4 and 5, we can evaluate the true operational influence of the cost function on the Conv-1D Autoencoder's latent space.

#### 1. Invariance of the Generated Pseudospectra

The visual and mathematical comparison of the generated pseudospectra across all configurations demonstrates a remarkable morphological invariance:

- **Preservation of Key Spectral Features:** Across all plots, the autoencoder consistently reconstructs the same four fundamental spectral typologies. The relative intensities, widths, and positions of the characteristic Raman D and G bands, as well as the underlying slopes of the uncorrected fluorescence baseline, remain virtually identical.
- **Consistency of Cluster Archetypes:** For instance, the highly resolved, low-baseline profile containing a sharp valley between the peaks (visible in Cluster 3 or Cluster 0 of the Shannon models and Cluster 2 of the  $\kappa = 0.5$  or Cluster 1 of the  $\kappa = 2.0$  Kaniadakis models) preserves its exact geometric identity. The network reconstructs the same physical features regardless of whether standard Euclidean distance or deformed metrics are applied.

#### 2. Alignment with Matrix Contingency Data

This spectral invariance perfectly mirrors the observations made on the cross-classification matrices in Figure 1. Because the clusters are defined by the broad, macroscopic shapes of the raw spectra—dominated by the prominent background luminescence—the underlying grouping criteria do not change when shifting from  $\kappa = 0.5$  to  $\kappa = 2.0$ . The minor sample reshuffling observed in the matrices represents subtle boundary adjustments rather than a meaningful reorganization of the latent taxonomy.

## Conclusion

These empirical results confirm that within a 1D Convolutional architecture, the choice of the statistical metric has a negligible effect on both the unsupervised partitioning and the generative synthesis of raw Raman signals. The spatial feature-extraction capacity of the convolutional filters shows a robust structural invariance. Consequently, for tasks involving raw carbonaceous material sorting and library generation, the classical Shannon-MSE and the Kaniadakis kappa-deformed loss functions can be considered functionally equivalent.

## Architectural Contrast: Conv-1D Robustness vs. Dense Models

The structural invariance observed across the clustering matrices and the generated profiles carries a crucial architectural significance. Rather than representing a mere lack of sensitivity, this "negative result" provides empirical proof that **1D Convolutional networks operate under a fundamentally different paradigm than Dense (Fully Connected) architectures** when mapping un-preprocessed spectral data.

In a traditional Dense network, every individual wavenumber or spectral bin is treated as an isolated, independent variable. Without spatial constraints, a Fully Connected layer is highly susceptible to shifts in the optimization metric geometry, as any modification to the loss function drastically alters how the weights balance global numerical distances.

Conversely, the Conv-1D layers extract features through localized spatial kernels. These convolutional filters capture structural correlations, local derivatives, and adjacent peak geometries (such as the definitive D and G molecular profiles) within sliding windows. Because the network prioritizes these local topological shapes over absolute, point-by-point Euclidean distances, the resulting latent space topology becomes highly robust against statistical deformations of the cost function.

Therefore, the morphological stability maintained from  $\kappa = 0.5$  to  $\kappa = 2.0$  highlights a major architectural divergence: while Dense models remain sensitive to the specific mathematical definition of the error space, Conv-1D architectures exhibit an intrinsic spatial invariance, anchoring the unsupervised taxonomy to the core structural features of the Raman signal regardless of the  $\kappa$ -deformation.

## References

Kaniadakis, G. (2001). Non-linear kinetics underlying generalized statistics. *Physica A: Statistical mechanics and its applications*, 296(3-4), 405-425.

Kaniadakis, G. (2002). Statistical mechanics in the context of special relativity. *Physical review E*, 66(5), 056125.

Sparavigna, A. C., & Gemini (Modello Linguistico di Google). (2025). Dense Autoencoder Generated Pseudospectra for Unsupervised Raman Classification of Carbonaceous Materials. Zenodo. <https://doi.org/10.5281/zenodo.16935868>

Sparavigna, A. C., & Gemini (Modello Linguistico di Google). (2025). Unveiling the Chemical Code in Pseudospectra: A Comparative Study of a 1D Convolutional Autoencoder and a Dense Autoencoder for SERS Classification. Zenodo. <https://doi.org/10.5281/zenodo.16912956>

Sparavigna, A. C., & Gemini (Modello Linguistico di Google). (2026). Beyond the Luminescence Trap: A Relativistic kappa-Dense Autoencoder for the Objective Taxonomy of Native Raman Spectra. Zenodo. <https://doi.org/10.5281/zenodo.20432996>

Sparkes, R. B., Maher, M., Blewett, J., Doğrul Selver, A., Gustafsson, Ö., Semiletov, I. P., & Van Dongen, B. E. (2018). Carbonaceous material export from Siberian permafrost tracked across the Arctic Shelf using Raman spectroscopy. *The Cryosphere*, 12(10), 3293-3309.

Sparkes, R. B., Maher, M., & Blewett, J. (2018) Raw data for paper title "Carbonaceous material export from Siberian permafrost tracked across the Arctic Shelf using Raman spectroscopy". [Dataset]. <https://doi.org/10.23634/MMUDR.00620205>