

Lawrence Berkeley National Laboratory

LBL Publications

Title

An interactive machine learning platform for analyzing multi-particle coincidence data from cold target recoil ion momentum spectroscopy

Permalink

<https://escholarship.org/uc/item/7976980f>

Journal

Review of Scientific Instruments, 97(5)

ISSN

0034-6748

Authors

Daoud, Hazem
Kumar, Sarvesh
Qian, Jin
[et al.](#)

Publication Date

2026-05-01

DOI

10.1063/5.0313735

Peer reviewed

SCULPT: An Interactive Machine Learning Platform for Analyzing Multi-Particle Coincidence Data from Cold Target Recoil Ion Momentum Spectroscopy

Hazem Daoud,^{*} Sarvesh Kumar, Jin Qian, Daniel Slaughter, and Thorsten Weber[†]

*Chemical Sciences Division, Atomic, Molecular and Optical Physics,
Lawrence Berkeley National Laboratory, 1 Cyclotron Road, Berkeley, California 94720, USA*

Tanny Chavez[‡]

*Advanced Light Source, Lawrence Berkeley National Laboratory,
1 Cyclotron Road, Berkeley, California 94720, USA*

(Dated: June 3, 2026)

We present SCULPT (**S**upervised **C**lustering and **U**ncovering **L**atent **P**atterns with **T**raining), a comprehensive software platform for analyzing tabulated high-dimensional multi-particle coincidence data from Cold Target Recoil Ion Momentum Spectroscopy (COLTRIMS) experiments. The software addresses critical challenges in modern momentum spectroscopy by integrating advanced machine learning techniques with physics-informed analysis in an interactive web-based environment. SCULPT implements Uniform Manifold Approximation and Projection (UMAP) for non-linear dimensionality reduction to reveal correlations in highly dimensional data. We also discuss potential extensions to deep autoencoders for feature learning, and genetic programming for automated discovery of physically meaningful observables. A novel adaptive confidence scoring system provides quantitative reliability assessments by evaluating user-selected clustering quality metrics with predefined weights that reflect each metric’s robustness. The platform features configurable molecular profiles for different experimental systems, interactive visualization with selection tools, and comprehensive data filtering capabilities. Utilizing a subset of SCULPT’s capabilities, we analyze photo double ionization data measured using the COLTRIMS method for 3-body dissociation of the D₂O molecule, revealing distinct fragmentation channels and their correlations with physics parameters. The software’s modular architecture and web-based implementation make it accessible to the broader atomic and molecular physics community, significantly reducing the time required for complex multi-dimensional analyses. This opens the door to finding and isolating rare events exhibiting non-linear correlations on the fly during experimental measurements, which can help steer exploration and improve the efficiency of experiments.

I. INTRODUCTION

Cold Target Recoil Ion Momentum Spectroscopy (COLTRIMS), often referred to as “reaction microscopy”, has revolutionized our ability to study atomic and molecular dynamics by enabling kinematically complete measurements of multi-particle fragmentation processes.[1–3] COLTRIMS allows the measurement of the three-dimensional momentum vectors of all ions and electrons produced in fundamental ionizing reactions in gaseous atoms, molecules, and complexes, providing unprecedented insight into quantum mechanical processes, correlations, and coherences.[4, 5]

The comprehensive nature of COLTRIMS measurements, while powerful, presents significant data analysis challenges. For instance, a typical experiment studying the photo double ionization of water resulting in three-body dissociation ($\text{H}^+ + \text{O}^+ + \text{H} + 2\text{e}^-$ or $\text{H}^+ + \text{H}^+ + \text{O} + 2\text{e}^-$), in which the 3D momentum of the neutral fragment was determined via momentum conservation, produces events characterized by 15 momentum

components (p_x , p_y , p_z for each of the five particles in each reaction), and the mass-to-charge ratio of each ion. When combined with derived physical quantities such as kinetic energy release (KER), relative angles between particles, electron energy sharing, and others, the effective dimensionality can easily exceed 50 features per event. Modern COLTRIMS experiments routinely generate datasets containing tens of millions of such high-dimensional events. These events are represented as rows in the list-mode file format, which records the events on a shot-by-shot basis. The features (observables and derived quantities) are stored in the columns of this large table.[6]

The current well-established analysis approaches for COLTRIMS data have significant limitations:

1. **Projection-based methods:** Researchers typically create one- or two-dimensional histograms of selected variables, potentially missing important multi-dimensional correlations that reveal reaction mechanisms.[7]
2. **Sequential filtering:** Applying cuts on selected parameters can introduce bias and obscure unexpected patterns in the data.[8]
3. **Manual feature engineering:** The analysis relies heavily on domain expertise to construct rele-

^{*} HDaoud@lbl.gov

[†] TWeber@lbl.gov

[‡] TanChavez@lbl.gov

vant observables, potentially missing non-intuitive relationships between observables or parameters.[9]

Recent advances in machine learning,[10] particularly in dimensionality reduction and unsupervised clustering, offer promising solutions to these challenges.[11] The application of machine learning to high-dimensional experimental data analysis has seen remarkable growth across multiple domains of experimental physics, chemistry, biology and engineering.[12, 13] In X-ray scattering and spectroscopy, variational autoencoders and other dimensionality reduction techniques have been successfully applied to learn latent representations of complex scattering patterns,[12] enabling researchers to identify similar structures through latent space clustering and to generate new structures through sampling and exploration of the learned manifold. Generative adversarial networks have been successfully applied to analyze ultrafast electron diffraction images, demonstrating the utility of deep learning approaches for extracting structural dynamics from high-dimensional experimental data where traditional analysis methods face ambiguity.[14] For small-angle X-ray scattering (SAXS) data, variational autoencoders have been employed to visualize large datasets in low-dimensional latent spaces, enabling rapid capture of key features such as similarity among scattering patterns and structural evolution trends.[15] In protein structure analysis, autoencoders trained to compress 3D shape information into 200-dimensional latent spaces have been combined with genetic algorithms to build 3D models consistent with scattering data,[16] demonstrating the power of latent space optimization for inverse problems.

Recent efforts have focused on creating integrated web-based platforms for ML-driven analysis of experimental data at scientific user facilities. The MLEExchange platform introduced a novel labeling pipeline that accelerates annotation of large scientific datasets using AI-guided tagging techniques, with interconnected graphical user interfaces for data reduction, classification, latent space exploration, and label assignment.[17] This approach has proven instrumental for pattern recognition in X-ray scattering data, enabling scientists to label large datasets efficiently while maintaining the ability to train and fine-tune customizable ML models within the same pipeline. Mass spectrometry imaging, which shares with COLTRIMS the challenge of high-dimensional, multi-parameter datasets, has benefited from variational autoencoder approaches that learn nonlinear spectral manifolds to reveal biologically relevant clusters.[18] Convolutional autoencoders have proven particularly effective at aggregating features and preserving low-abundant signals in their latent space representations,[19] addressing challenges similar to those faced in identifying rare fragmentation channels in COLTRIMS data.

The Uniform Manifold Approximation and Projection (UMAP) technique[20] has emerged as a particularly powerful tool for dimensionality reduction in biological and physical sciences.[21, 22] Compared to other dimensionality reduction methods, UMAP pro-

vides high reproducibility, and meaningful organization of clusters,[21] while preserving both local neighbor relations and aspects of global structure. Applications in molecular dynamics simulations of biomacromolecules have demonstrated UMAP’s superior performance when compared with linear reduction methods and competitive performance with other nonlinear techniques.[23] In particle physics, machine learning has transformed data analysis and simulation, with applications ranging from boosted decision trees for classification to various types of neural networks for pattern recognition and event reconstruction.[24, 25] Reinforcement learning approaches have been applied to hierarchical clustering problems in particle physics, demonstrating that such tasks can be phrased as Markov Decision Processes.[26]

However, existing machine learning tools are typically not tailored to the specific requirements of tabulated momentum spectroscopy data, lacking physics-informed constraints, configurable molecular systems for different experimental targets, adaptive quality assessment for clustering reliability, and interactive exploration tools for multi-dimensional correlations. In parallel with the present work, Venkatachalam et al. [27] recently reported on ML-assisted Coulomb explosion imaging (CEI) guided by classical simulations to extract molecular geometries and distinguish between isomers of mid-sized molecules from list-mode data. In another concurrent study, Li et al. [28] employed a deep generative neural network to infer molecular structures from ion momentum distributions generated during the rapid Coulomb explosion of molecules with up to nine atoms. In both approaches, the complete CEI, where all atomic ionic fragments of the molecule are detected, is currently a prerequisite, and based on this, the authors demonstrate its potential to investigate molecular transformation in real-time with unprecedented detail. Beyond complete CEI, which is well suited for strong-field laser and XFEL CEI experiments, the momentum-spectroscopy community will benefit from an integrated platform that unifies modern dimensionality-reduction methods with domain-specific analysis tools in an accessible, web-based environment capable of correlating electrons, ions, and at least one neutral fragment in a general and flexible way.

We present SCULPT, an integrated software platform that bridges this gap by combining state-of-the-art machine learning techniques with physics-informed analysis tools specifically designed for multi-particle coincidence list-mode data. The key innovations of SCULPT include:

1. **Physics-based dimensionality reduction:** The integration of UMAP[20] with automated physics feature calculation ensures that dimensionality reduction preserves physically meaningful structures.
2. **Adaptive and quantitative reliability assessment:** A confidence scoring system that evaluates user-selected clustering quality metrics through weighted combination provides interpretable reliability assessments of UMAP results.

3. **Interactive exploration of multidimensional data by arbitrary dimensionality reduction:** Dynamic visualization with selection and filtering tools enables a hypothesis-driven analysis unrestricted by conventional 1-D or few-dimensional projections.
4. **Flexible configurable molecular systems:** A profile-based system for defining arbitrary molecular configurations calculates appropriate mass- and charge-dependent features.
5. **Automated feature discovery:** Genetic programming algorithms[29] that discover mathematical combinations of input features aid in discovering relevant clusters and correlations.

II. SOFTWARE ARCHITECTURE AND IMPLEMENTATION

SCULPT is implemented as a web-based application using Python and Plotly Dash framework,[30] chosen for its ability to create reactive, interactive visualizations without requiring client-side programming. The platform follows a modular design that integrates data processing, machine learning analysis, quality assessment, and interactive visualization capabilities into a cohesive workflow for COLTRIMS data exploration.

A. Data Processing and Physics Feature Calculation

The data processing layer handles importing COLTRIMS data files and implements a novel molecular configuration system. Users can define molecular profiles by specifying particle types, masses, and charges (see Fig. 1).

Moreover, it enables accurate physics calculations for various molecular photoionization experiments, with targets such as different isotopologues (e.g., H_2O , D_2O , HDO), as well as other molecular systems, without source code modification.

The physics feature calculator automatically computes:

- Individual particle kinetic energies: $E_i = p_i^2/(2m_i)$
- Kinetic Energy Release of dissociation processes: $\text{KER} = \sum E_{\text{ions}}$
- Sum of electron energies: $\text{EESum} = \sum E_{\text{electrons}}$
- Relative angles between all particle pairs
- Energy sharing parameters such as electron energy sharing: $\text{EESharing} = E_{\text{electron}}/\sum E_{\text{electrons}}$

B. Machine Learning Components

1. UMAP Implementation

SCULPT employs UMAP[20] for non-linear dimensionality reduction of high-dimensional COLTRIMS data. UMAP constructs a topological representation of the high-dimensional data by modeling each data point's relationship to its nearest min, then optimizes a low-dimensional embedding that preserves these local neighborhood structures while maintaining global data topology. This approach is particularly effective for COLTRIMS data as it reveals non-linear correlations between momentum components and derived physics quantities that are not apparent in traditional linear projection methods. The user selects a set of calculated physics features (kinetic energies, angular distributions, correlation parameters) or original momentum components as the inputs to be embedded into the UMAP algorithm. The implementation preserves file labels throughout the dimensionality reduction process, enabling direct comparison of multiple datasets or experimental conditions within the same 2D visualization space. Key UMAP parameters including the number of neighbors ($n_{\text{neighbors}}$, controlling local versus global structure preservation) and minimum distance (min_dist , controlling cluster compactness) are user-configurable, with typical values of $n_{\text{neighbors}} = 15$ and $\text{min_dist} = 0.1$ providing effective visualizations for COLTRIMS datasets.

2. Deep Autoencoder Architecture

The autoencoder module implements a symmetric neural network architecture using PyTorch for feature learning from high-dimensional COLTRIMS data. The network consists of an encoder that progressively compresses the input feature space through multiple hidden layers to a low-dimensional latent representation, and a symmetric decoder that reconstructs the original features from this compressed representation. The latent space dimensionality (typically 2-10 dimensions) is user-configurable, with the compression forcing the network to learn the most salient features that capture the essential variance in the fragmentation dynamics. Training employs mean squared error reconstruction loss with the Adam optimizer, and the resulting latent features can be visualized directly or used as input to subsequent UMAP projection and clustering analysis. This approach is particularly effective for discovering non-linear correlations between momentum components and derived physics quantities that may not be apparent through traditional feature selection methods.

SCULPT Molecular Configuration and Data Processing Workflow




FIG. 1: Flowchart showing the molecular configuration and data processing workflow.

3. Genetic Programming for Feature Discovery

SCULPT implements genetic programming using symbolic regression to discover interpretable feature combinations.[29, 31] Genetic programming has proven effective for automated feature construction in scientific data analysis, particularly when the goal is to discover non-trivial mathematical relationships between input variables that may not be apparent through traditional feature engineering.[32, 33] This approach is explicitly chosen to discover non-trivial features that might provide superior cluster separation compared to standard physics parameters (like energy and momentum conservation). The method has been successfully applied to experimental physics data, where dimensional consistency and physical interpretability of constructed features are critical requirements.[31] By evolving pop-

ulations of mathematical expressions through selection, crossover, and mutation operations, the genetic programming module can identify complex combinations of momentum components, energies, and angular features that capture underlying correlations in fragmentation patterns, thereby enhancing the ability to distinguish between different quantum states or reaction pathways in COLTRIMS data.

C. Adaptive Quality Assessment

A key innovation in SCULPT is the adaptive confidence scoring system that provides quantitative reliability assessments for clustering results. The system evaluates user-selected clustering quality metrics, each automatically assigned predefined weights and reliability

scores based on their robustness and general applicability. Users select which metrics to calculate through the interface, allowing flexibility and balance between computational cost and analysis depth. The confidence score is then computed as a weighted combination of only the selected metrics, with weights normalized to account for that subset of metrics. Metrics are organized into three tiers based on their general reliability: Tier 1 metrics (silhouette score, Hopkins statistic) are highly reliable and recommended for all analyses; Tier 2 metrics (cluster stability, physics consistency, Calinski-Harabasz index) provide additional validation with moderate reliability; and Tier 3 metrics (Davies-Bouldin index) are included with heavily reduced weight due to known limitations. The only automatic exclusion occurs for the physics consistency metric, which is omitted when physics features are unavailable in the dataset. These reliability tiers and associated weights were empirically determined through validation on the D₂O double ionization dataset with known ground truth quantum state assignments.[34] The tier assignments reflect performance patterns observed when separating the eight distinct dication states in this well-characterized system, where the metrics' ability to correctly identify physically meaningful clusters could be verified against established quantum state separations. **The adaptive confidence score is intended as a heuristic guide for exploratory analysis rather than a rigorous statistical measure of physical correctness. In practice, scores are categorized as follows: ≥ 0.65 (High) indicates reliable cluster separation suitable for drawing preliminary conclusions; $0.50\text{--}0.65$ (Moderate) suggests reasonably reliable results that warrant further validation; $0.35\text{--}0.50$ (Low) indicates results should be used with caution and may benefit from alternative feature selection or UMAP parameters; and ≤ 0.35 (Very Low) suggests the clustering may not capture meaningful structure.**

Tier 1 metrics (highest reliability, recommended for all analyses):

1. **Silhouette Score** $S \in [-1, 1]$: This metric measures both cluster cohesion (how similar a point is to others in its own cluster) and separation (how dissimilar it is to points in neighboring clusters).[35] For each data point i , the silhouette coefficient is calculated as $s(i) = \frac{b(i) - a(i)}{\max\{a(i), b(i)\}}$, where $a(i)$ is the mean distance to other points in the same cluster and $b(i)$ is the mean distance to points in the nearest neighboring cluster. The overall silhouette score is the mean across all points. Values near +1 indicate well-separated clusters, values near 0 suggest overlapping clusters, and negative values indicate potential misclassification. This metric receives a weight of $w = 0.35$ (contributing 35% to the final confidence score) and a reliability score of $r = 0.9$ (indicating 90% confidence in its assessment validity). The weight w is the contribution to the overall confidence score, and r represents the

reliability or confidence in the validity of the metric. The weights sum to unity across all active metrics, and reliability scores modulate the influence of each metric based on empirical performance across diverse clustering scenarios.

2. **Hopkins Statistic** $H \in [0, 1]$: This metric assesses the clustering tendency of the dataset by measuring the probability that the data is generated from a uniform distribution versus containing meaningful clusters.[36, 37] The statistic compares the nearest-neighbor distances for a sample of real data points against those for randomly generated points within the same space. Values of $H > 0.75$ indicate strong clusterability (the data significantly deviates from uniform randomness), while $H < 0.5$ suggests the data may be uniformly distributed with no inherent cluster structure. This metric receives $w = 0.25$ and $r = 0.85$, reflecting its importance in determining whether clustering is appropriate for the dataset. Here, H refers specifically to the Hopkins statistic value.

Tier 2 metrics (moderate reliability, useful for comprehensive assessment):

1. **Cluster Stability** $\in [0, 1]$: Measures the reproducibility of cluster assignments under small perturbations by injecting Gaussian noise (5% of feature standard deviation) into the data and comparing the resulting cluster assignments with the original using the Adjusted Rand Index.[38] High stability values (> 0.8) indicate robust clusters that are insensitive to minor variations in the data, while low values suggest the clustering may be capturing noise rather than genuine structure. This metric receives $w = 0.15$ and $r = 0.7$.
2. **Physics Consistency** $\in [0, 1]$: This metric is automatically excluded when physics features are unavailable in the dataset (e.g., when analyzing autoencoder latent representations). It is a domain-specific validation metric that measures the ratio of within-cluster variance to between-cluster variance for key physics parameters. The metric evaluates a predefined set of physics quantities: KER, EESum, EESharing, individual particle energies (ions and electrons), and total energy. For each available physics parameter, the metric calculates $PC_i = \frac{\text{Var}_{\text{between}}}{\text{Var}_{\text{between}} + \text{Var}_{\text{within}}}$, where $\text{Var}_{\text{between}}$ is the variance of cluster means weighted by cluster size, and $\text{Var}_{\text{within}}$ is the weighted average of within-cluster variances. The overall physics consistency is the mean across all evaluated parameters. Values near 1 indicate that clusters correspond to physically distinct states with homogeneous physics parameter distributions within each cluster, even if these physics parameters were not explicitly used as clustering features. This provides an independent validation that the clustering captures mean-

ingful physical differences rather than arbitrary data structure. This metric receives $w = 0.2$ and $r = 0.8$.

3. **Calinski-Harabasz Index** $\in [0, \infty)$: Measures cluster separation by computing the ratio of between-cluster dispersion to within-cluster dispersion in the UMAP latent space.[39] The index is defined as $CH = \frac{\text{tr}(B_k)}{\text{tr}(W_k)} \times \frac{N-k}{k-1}$, where B_k is the between-cluster dispersion matrix, W_k is the within-cluster dispersion matrix, N is the number of data points, and k is the number of clusters. Higher values indicate better-defined clusters with greater separation in the reduced-dimensional space. This metric receives $w = 0.1$ and $r = 0.6$.

Tier 3 metrics (lower reliability, heavily down-weighted):

1. **Davies-Bouldin Index** $\in [0, \infty)$: Measures the average similarity ratio between each cluster and its most similar neighbor, where similarity is defined as the ratio of within-cluster scatter to between-cluster separation.[40] Lower values indicate better clustering. However, this metric is known to be sensitive to cluster shape assumptions (favoring spherical clusters) and can produce misleading results for elongated or irregular cluster geometries common in UMAP embeddings.[41] Due to these limitations, it receives minimal weight ($w = 0.005$) and low reliability ($r = 0.4$).

Additionally, the system tracks the **Noise Ratio** $\in [0, 1]$, representing the fraction of points classified as noise by the clustering algorithm. SCULPT implements DBSCAN (Density-Based Spatial Clustering of Applications with Noise)[42] as the primary clustering algorithm, which automatically identifies noise points as those not belonging to any dense region. For the main UMAP analysis and custom scatter plots, DBSCAN uses automatic parameter optimization where the algorithm searches for the optimal epsilon value (distance threshold) while maintaining a fixed minimum sample requirement of 5 neighbors. The optimization procedure tests epsilon values in the range $[0.1, 1.0]$ and selects the configuration that maximizes the number of clusters while keeping the noise ratio below 50%. In the advanced genetic programming module, users have full control over both the epsilon and minimum samples parameters, allowing for fine-tuning based on specific dataset characteristics. Points are labeled as noise if they have fewer than the minimum number of neighbors within the specified radius. The noise ratio influences both reliability assessment and performance bonuses, with very low noise ratios (< 0.05) indicating clean cluster separation.

The adaptive confidence score C serves as a summary metric that combines all individual quality assessments into a single interpretable value. It employs a reliability-weighted calculation:

$$C = \frac{\sum_i w_i \cdot r_i \cdot n_i}{\sum_i w_i \cdot r_i}, \quad (1)$$

where w_i is the weight for metric i , r_i is its reliability score, and n_i is the normalized value of metric i . The sum ranges only over the metrics selected by the user for calculation. The normalization function n_i transforms each metric to a $[0, 1]$ scale using metric-specific scaling functions designed to provide realistic assessments across the full range of possible values.

The system incorporates several advanced features to ensure robust assessment:

Adaptive normalization employs metric-specific scaling functions that account for the different value ranges and distributions of each metric. For the silhouette score, the transformation is $n_S = 0.5(S + 1)$, mapping $[-1, 1]$ to $[0, 1]$. For the Hopkins statistic, $n_H = H$ (already in $[0, 1]$). For the Calinski-Harabasz index, a logarithmic transformation $n_{CH} = \min(1, \log(1 + CH)/\log(1000))$ prevents extremely large values from dominating. For the Davies-Bouldin index, $n_{DB} = \max(0, 1 - DB/3)$ transforms lower-is-better to higher-is-better with appropriate scaling. These functions avoid overly harsh penalties for moderate clustering while preserving sensitivity to exceptional results, with normalized values capped at 0.98 to reserve perfect scores for theoretical limits.

Critical threshold management applies contextual penalties to prevent overconfident assessments of poor clusterings. Severe failures (silhouette score $S < -0.1$ or Hopkins statistic $H < 0.3$) cap the confidence at 0.4, indicating fundamental clustering problems. Borderline cases are identified as those with $0.2 < S < 0.4$ or $0.5 < H < 0.7$, which receive moderate penalties (confidence capped at ≤ 0.7) to reflect ambiguous cluster quality.

Performance bonuses are applied using asymptotic scaling to reward exceptional clustering without inflating confidence unrealistically. Four bonus conditions are evaluated: exceptional silhouette scores ($S > 0.6$, +0.1 bonus), very low noise ratios (noise < 0.05 , +0.05 bonus), high stability (stability > 0.8 , +0.05 bonus), and strong clustering tendency ($H > 0.8$, +0.05 bonus). The bonus system uses asymptotic scaling: $C_{\text{new}} = C + (0.95 - C) \cdot \frac{\text{bonus}}{0.95}$ to maintain realistic confidence bounds and prevent scores from approaching 1.0.

Uncertainty quantification provides confidence intervals around the final score to communicate the reliability of the assessment. The base uncertainty is calculated as $\sigma_{\text{base}} = 0.1 + \max(0, (3 - n_{\text{metrics}})) \times 0.05$, where n_{metrics} is the number of active metrics used in the calculation. This function was empirically derived through iterative testing on the eight known dication states in the D₂O double ionization dataset,[34] where the uncertainty model parameters (0.1 baseline and 0.05 penalty per missing metric) were tuned to reflect observed variations in metric reliability. The 0.1 baseline represents the estimated inherent measurement un-

certainty, while the penalty term increases uncertainty when fewer metrics are available (reflecting reduced information). The base uncertainty is further adjusted for the average reliability of active metrics: $\sigma_{\text{total}} = \sigma_{\text{base}} + \max(0, (0.8 - r_{\text{avg}})) \times 0.1$, where $r_{\text{avg}} = \frac{\sum_i r_i}{n_{\text{metrics}}}$. This adjustment increases uncertainty when relying on less reliable metrics.

D. Advanced Data Filtering and Selection

SCULPT incorporates sophisticated data filtering capabilities that enable users to refine datasets at multiple stages of the analysis pipeline, enhancing both computational efficiency and analytical focus. The system provides three primary filtering mechanisms: (1) feature-based selection, (2) density-based filtering, and (3) physics parameter filtering.

1. Feature Selection Framework

The feature selection system organizes available variables into hierarchical categories based on their physical significance and computational origin. Features are automatically categorized as follows:

Original Momentum Components: Raw momentum measurements (p_x, p_y, p_z) for each detected particle, providing the fundamental kinematic information from experimental data.

Momentum Magnitudes: Derived scalar quantities $|\vec{p}_i| = \sqrt{p_{x,i}^2 + p_{y,i}^2 + p_{z,i}^2}$ for each particle, offering magnitude-based clustering without directional bias.

Energy Variables: Comprehensive energy features including individual particle kinetic energies, Kinetic Energy Release representing the sum of ion energies, electron energy sum, electron energy sharing ratios, and total system energy.

Angular Features: Laboratory-frame polar angles $\theta_i = \arccos(p_{z,i}/|\vec{p}_i|)$ and azimuthal angles $\phi_i = \arctan 2(p_{y,i}, p_{x,i})$ for each particle, enabling analysis of angular distributions and correlations.

Correlation Features: Inter-particle relationships including dot products $\vec{p}_i \cdot \vec{p}_j$, momentum differences, and angles between particle pairs $\cos(\alpha_{ij}) = \frac{\vec{p}_i \cdot \vec{p}_j}{|\vec{p}_i||\vec{p}_j|}$.

The system provides hierarchical selection controls with category-level selection and individual feature toggling, which enables efficient exploration of the high-dimensional feature spaces, while maintaining the interpretability of the selected subset.

2. Density-Based Filtering

Density-based filtering addresses the challenge of computational scalability and noise reduction by selectively retaining high-density regions in both UMAP projections

and direct feature scatter plots. The implementation uses efficient grid-based density estimation that operates on any two-dimensional data representation:

$$\rho(x_i, y_i) = \frac{1}{N} \sum_{j=1}^N G_\sigma(x_i - x_j, y_i - y_j), \quad (2)$$

where G_σ is a Gaussian kernel with bandwidth parameter σ (user-adjustable from 0.01 to 1.0), and (x_i, y_i) represents coordinates in either the 2D UMAP embedding space or direct feature space (e.g., energy vs. angle scatter plots). The system creates a 100×100 grid histogram of the selected coordinates, applies Gaussian smoothing with $\sigma_{\text{smooth}} = 10\sigma$, and assigns density values to each point based on grid interpolation.

Points are filtered using a percentile-based threshold τ_p , where users specify the percentile (0-100%) of density values to retain. Points with density $\rho \geq \tau_p$ are preserved, effectively removing sparse outliers while maintaining the core cluster structure. This approach significantly reduces computational overhead for subsequent analyses, while preserving essential clustering patterns.

For scatter plot applications, density filtering enables users to focus on statistically significant regions within the original feature space by removing experimental noise and low-statistics outliers that might otherwise obscure underlying physical correlations. The same mathematical framework applies whether filtering UMAP embeddings or direct two-dimensional projections of calculated physics parameters.

3. Physics Parameter Filtering

Physics parameter filtering enables domain-specific data selection based on physically meaningful quantities. The system supports filtering on commonly used physics parameters including kinetic energy release (KER), individual particle kinetic energies for ions and electrons, electron energy sums and energy sharing ratios, inter-particle angular correlations, and momentum differences between particles. This filtering capability enables selection of specific energy regimes, geometric configurations, or correlation patterns corresponding to different physical processes.

For each selected parameter, the system automatically determines the data range and provides an interactive range slider with intelligent binning. The parameter range is dynamically calculated as $[\lfloor \min(P) \rfloor, \lceil \max(P) \rceil]$ where P represents the selected parameter values across all loaded datasets. The filtering condition retains data points where $P_{\min} \leq P_i \leq P_{\max}$, based on user-specified bounds.

4. Integrated Filtering Pipeline

The filtering system operates through an integrated pipeline where filters can be applied sequentially:

1. **Feature Selection:** Users first select relevant features for dimensional reduction, reducing computational complexity and focusing on scientifically relevant dimensions.

2. **UMAP Projection:** The selected features undergo UMAP embedding to create the 2D visualization space.

3. **Density Filtering:** Optional density-based filtering removes low-density outliers from the UMAP space, improving signal-to-noise ratio.

4. **Physics Filtering:** Optional parameter-based filtering selects physically meaningful subsets for targeted analysis.

Each filtering step preserves data provenance, maintaining links between filtered subsets and original experimental data. The system provides detailed filtering statistics, which includes retention percentages and per-file event counts, enabling users to assess the impact of each filtering operation.

This multi-stage filtering approach enables efficient exploration of large experimental datasets while maintaining physical interpretability and statistical rigor, supporting both exploratory data analysis and hypothesis-driven investigation of specific reaction mechanisms.

E. Interactive Visualization and Selection Tools

SCULPT provides multiple visualization modes with consistent interaction patterns:

1. **Scatter plots:** Interactive 2D projections with lasso and box selection tools
2. **Density heatmaps:** Kernel density estimation for identifying high-density regions

Selection tools are implemented using Plotly’s built-in selection capabilities. These tools empower the user to identify and isolate groups of events and save them in separate files for further processing. To further understand the physics that correlates events in any one cluster, SCULPT enables the user to visualize events in scatter plots or heat maps according to any of the original features such as total energy, kinetic energy release (KER), etc. This allows for vetting the data, visualizing the clusters identified by SCULPT on lower-dimensional phase spaces and understanding the physics behind the clustering.

F. Design Philosophy and Default Parameters

In this section we shed light on the default parameter

choices implemented in SCULPT.

Confidence score weighting: The tiered metric system emphasizes local density structure (silhouette score, weight 0.35) and intrinsic clusterability (Hopkins statistic, weight 0.25) because these properties are most relevant for identifying physically distinct fragmentation channels in momentum space. Metrics known to be sensitive to cluster shape assumptions, such as the Davies-Bouldin index, are heavily downweighted (weight 0.005) since UMAP embeddings frequently produce non-spherical cluster geometries. This weighting scheme may undervalue clustering quality in datasets where global structure is more important than local density, or where clusters have highly uniform, spherical shapes.

DBSCAN parameter selection: The automatic epsilon optimization maximizes the number of detected clusters subject to a 50% noise threshold. This heuristic favors segmentation—revealing fine structure that might otherwise be merged—over conservative clustering that minimizes false positives. For applications requiring high-purity cluster assignments, users should consider manually adjusting epsilon toward larger values or using the advanced analysis module where full control over DBSCAN parameters is available.

UMAP parameters: The default values ($n_{\text{neighbors}} = 15$, $\text{min_dist} = 0.1$) balance local and global structure preservation for typical COLTRIMS dataset sizes (10^4 – 10^6 events). Smaller $n_{\text{neighbors}}$ values emphasize local structure and may reveal finer sub-clustering, while larger values preserve more global topology at the cost of local detail. These parameters should be adjusted based on dataset size and the scale of structure being investigated.

The current defaults may perform poorly for datasets with highly imbalanced cluster sizes, where small but physically important clusters may be classified as noise; for data with continuous gradients rather than discrete clusters; or for very small datasets (< 1000 events) where the Hopkins statistic and stability metrics become unreliable. In such cases, users should rely more heavily on visual inspection and physics-based validation rather than the confidence score alone. Additionally, the filtering features incorporated in SCULPT serve to aid users identify clusters that would otherwise be challenging to isolate.

III. USER INTERFACE AND WORKFLOW

A. Interface Organization

SCULPT’s interface is organized into five main sections (see Fig. 2):

1. **Data & Configuration:** File upload and molecular profile management
2. **Basic Analysis:** Probabilistic machine learning (UMAP embedding) and custom feature plots

3. **Selection & Filtering:** Interactive data selection and filtering tools
4. **Advanced Analysis:** Re-analysis of selected subsets
5. **Machine Learning:** Autoencoder and genetic programming modules

B. Typical Analysis Workflow

A typical analysis workflow proceeds as follows:

1. **Data Import:** Upload COLTRIMS data files and select appropriate molecular configuration
2. **Initial Exploration:** Run UMAP to visualize data structure.
3. **Cluster Identification:** Use DBSCAN with automatic parameter optimization or just rely on manual data separation by identifying clusters visually.
4. **Quality Assessment:** Review confidence scores and individual metrics.
5. **Refinement:** Select interesting regions for detailed analysis.
6. **Feature Discovery:** Optionally, apply genetic programming to attempt to discover new correlations.
7. **Validation:** Verify results by understanding the underlying physics, for example via plotting the identified clusters on physics-informed phase spaces.
8. **Export:** Save filtered datasets and discovered features as labeled list-mode data for further processing.

All these steps are to be reiterated with the data in saved files for further processing.

IV. CASE STUDY: D₂O DOUBLE IONIZATION

To demonstrate SCULPT’s capabilities with a well-conditioned example, we present a reanalysis of the D₂O double ionization following single photon absorption at 61 eV. Particularly, we considered the $D_2O^{2+} \rightarrow D^+ + D^+ + O$ dissociation channel that was previously analyzed for H₂O by Reedy et. al.[34] We created ground truth data based on the analysis in that work to isolate each of the quantum states of the water dication following dissociative photo- double ionization. This resulted in eight event list data files, containing the 3D momentum vector components of each particle for each event, for each quantum state: 3A_2 , 3B_1 , 3B_2 (producing oxygen O ³P), 1^1A_1 , 2^1A_1 , 1^1B_1 , 1^1B_2 (producing O ¹D), 3^1A_1 (producing O ¹S).

A. Experimental Data

The dataset, consisting of the eight aforementioned files, contains around 1,900,000 coincidence events where four of the five particles were detected: $D^+ + D^+ + O + e^- + e^-$. Each event is characterized by 15 momentum components (p_x, p_y, p_z for each of the five particles). ~~For computational efficiency, all of the UMAP analyses presented here were conducted on a random 1% sample of the full dataset. This sampling approach allows for rapid iterative exploration while preserving the essential clustering structure of the data. The sample size is chosen to provide reasonably low statistical uncertainties and sufficient density over the measured phase space to visualize clusters. After identifying clusters of interest on a sampled subset, we recommend users to re-run UMAP with different sampling fractions (ideally up to the full dataset) to verify that the cluster structure is preserved.~~

B. Analysis Results

~~To demonstrate SCULPT’s capabilities while maintaining computational efficiency, all UMAP analyses presented here were conducted on a random 1% sample of the full dataset. This sampling approach allows for rapid iterative exploration while preserving the essential clustering structure of the data.~~

1. UMAP Visualization

The initial UMAP analysis revealed five distinct clusters, separated by white space in the 2D projection, as shown in Fig. 3. This UMAP analysis was run by selecting the features KER, EESum, Total Energy (KER + EESum), and the angle between α_{12} between the ion 1 and ion 2 momenta. Two of the clusters consist of one quantum state exclusively, while the other three clusters include two or more quantum states. These clusters need to be further analyzed to isolate each of the quantum states contained in them. Typically, the user would attempt to separate quantum states within large clusters while being guided by the confidence score and the visual check of achieving clear separations in UMAP plots. As an example, in Fig. 4, we show the cluster containing the three states that was separated with the Lasso selection tool.

After isolating this cluster, we run UMAP on it again, this time while selecting the features KER, EESum, and Total Energy. The result, depicted in Fig. 5, shows that the group is now separated into two distinct clusters. One of the clusters contains the two states producing the oxygen atom in the ¹D state, specifically the water dication states 1^1B_2 (light-blue in Fig. 5) and 2^1A_1 (orange). The calculated confidence score is 0.70 confirming a high confidence in separating the data according to the selected features. We select the cluster containing the two states

SCULPT Interactive Analysis and Quality Assessment Workflow




FIG. 2: Flowchart showing the interactive analysis and quality assessment workflow.

and run UMAP analysis on it, again using the features KER, EESum, Total Energy, and α_{12} . The cluster is subsequently separated into two distinct clusters, each containing a single quantum state, as shown in Fig. 6. A high confidence score of 0.79 is calculated. Following the same general procedures described for the examples above, all eight quantum states in the present data are separable into individual clusters.

2. Cluster Identification and Physics Interpretation

A detailed analysis of the data enabled the isolation of the following clusters stemming from the $D_2O^{2+} \rightarrow D^+$

+ D^+ + O fragmentation channel:

Cluster 1 (dark-blue in Fig. 3, 2.6% of events): $D_2O^{2+} (^1B_1)$ dissociating to $D^+ + D^+ + O(^1D)$

- A singlet dication state that dissociates with a peak KER of ~ 4.3 eV and a β of $\sim 149^\circ$. The kinematic signatures of this state overlap with those of the triplet dication states 3B_1 and $^3^1A_1$, making clean separation more challenging than some other examples presented here.

Cluster 2 (red in Fig. 3, 3.0% of events): $D_2O^{2+} (^3B_1)$ dissociating to $D^+ + D^+ + O(^3P)$

- A triplet dication state with a peak KER of ~ 4.3




FIG. 3: UMAP projection of D_2O double ionization data showing five identified clusters. The calculated confidence score is 0.71. Colors represent different dication states and are for visualization purposes only.




FIG. 4: Selected cluster containing three dication states. Colors represent different dication states and are for visualization purposes only.

eV and a β of $\sim 146^\circ$. It is characterized by a narrow β distribution.

Cluster 3 (dark-green in Fig. 3, 4.6% of events): D_2O^{2+} (3^1A_1) dissociating to $D^+ + D^+ + O(^1S)$

- The only dication state leading to the $O(^1S)$ asymptote. It is a singlet state with a peak KER of ~ 11.4 eV and a β of $\sim 111^\circ$.

Cluster 4 (purple in Fig. 3, 8.2% of events): D_2O^{2+} (1^1A_2) dissociating to $D^+ + D^+ + O(^1D)$

- A singlet dication state with a peak KER of ~ 7.6 eV and a β of $\sim 126^\circ$.




FIG. 5: UMAP projection of D_2O double ionization data of the selected cluster. Separation into two clusters is evident. The bottom cluster contains data from two dication states. The calculated confidence score is 0.70. Colors represent different dication states and are for visualization purposes only.

Cluster 5 (orange in Figs 3 - 6, 10.7% of events): D_2O^{2+} (2^1A_1) dissociating to $D^+ + D^+ + O(^1D)$

- A singlet dication state characterized by a high β (protons nearly opposite directions) and a peak KER of ~ 7.7 eV.

Cluster 6 (light-blue in Figs 3 - 6, 18.2% of events): D_2O^{2+} (1^1B_2) dissociating to $D^+ + D^+ + O(^1D)$

- A singlet dication state with a peak KER of ~ 9.8




FIG. 6: UMAP projection of the lower cluster in Fig. 5. Separation of the data into two dication states is achieved. The calculated confidence score is 0.79. Colors represent different dication states and are for visualization purposes only.

eV and a β of $\sim 143^\circ$. It is a strong contributor to the $O(^1D)$ asymptote.

Cluster 7 (pink in Fig. 3, 19% of events): $D_2O^{2+} (^3A_2)$ dissociating to $D^+ + D^+ + O(^3P)$

- A triplet dication state with a peak KER of ~ 8 eV and a β of $\sim 121^\circ$.

Cluster 8 (light-green in Figs 3 - 6, 33.7% of events): $D_2O^{2+} (^3B_2)$ dissociating to $D^+ + D^+ + O(^3P)$

- A triplet dication state with a peak KER of ~ 9.7 eV and a β of $\sim 139^\circ$. It is a strong contributor to the $O(^3P)$ asymptote.

3. Quality Metrics

SCULPT's adaptive confidence scoring provided the following assessment for the initial UMAP analysis in Fig. 3. These values can vary slightly from run to run due to the random sampling approach.

- Overall confidence: 0.71 (High reliability)
- Silhouette score: 0.1324
- Hopkins statistic: 0.9769
- Stability: 0.9996
- Physics consistency: 0.3184
- Calinski Harabasz: 2338.4791

- Davies Bouldin: 0.7306

As clusters containing multiple quantum states are progressively separated through iterative analysis, the reliability score increases. For example, the confidence score for the UMAP embedding in Fig. 5 is 0.70, and when it is further sub-clustered, as shown in Fig. 6, the confidence score becomes 0.79. In contrast, selecting only KER, EESum and Total Energy as UMAP features results in poor separation, as shown in Fig. 7, and a poor score of 0.14. Hence, the confidence score, combined with visual inspection of cluster separation, guides the iterative selection of data subsets and appropriate features for subsequent clustering refinement.

It is important to note that individual metrics may exhibit low absolute values while the overall clustering remains physically meaningful. This is expected to be common in COLTRIMS data where quantum states exhibit inherent overlap in momentum space. For example, the silhouette score of 0.13 observed in Fig. 3 reflects the fact that several dication states share similar kinematic signatures (e.g., the 3B_1 and 1B_1 states both have peak KER ~ 4.3 eV), resulting in genuinely overlapping clusters in the UMAP projection. The confidence score accounts for this by weighting multiple complementary metrics—notably the Hopkins statistic (0.98 in Fig. 3), which indicates strong underlying clustering tendency despite the modest silhouette value. The iterative refinement workflow demonstrated in Section IV shows how progressively isolating clusters leads to improved confidence scores (0.70 $\rightarrow$ 0.79) as overlapping states are separated.




FIG. 7: UMAP projection of D_2O double ionization data showing poor clustering resulting from selecting only KER, EESum and Total Energy as UMAP features. The calculated confidence score is 0.14. Colors represent different dication states and are for visualization purposes only.

V. CONCLUSIONS

SCULPT represents a significant advance in the analysis of multi-particle coincidence data from momentum spectroscopy experiments. By combining modern machine learning techniques with physics-based parameters for dimensionality reduction, the software enables researchers to extract detailed information from high-dimensional data more efficiently than previously established analytical tools. We expect that SCULPT could also enable the identification and isolation of quantum states and dynamical mechanisms that could otherwise be left hidden in congested momentum spectra. The adaptive confidence scoring system provides quantitative reliability assessments, addressing a critical need for objective quality metrics in exploratory data analysis.

The case study of D₂O double ionization demonstrates SCULPT’s ability to identify subtle fragmentation channels, with several being diffuse or diluted by competing processes, and discover non-intuitive feature combinations that enhance cluster separation. The web-based implementation and modular architecture make the software accessible to the broader community while enabling customization for specific experimental requirements.

VI. FUTURE DIRECTIONS

We have not utilized the entire spectrum of capabilities and options of SCULPT in our analysis just yet. Near future studies on different datasets will explore other capabilities of SCULPT such as advanced data filtering and the implementation of deep autoencoders. As momentum spectroscopy techniques continue to be further developed, producing ever-larger and more complex datasets, tools like SCULPT will become essential for extracting physical insights.

The next development step will incorporate digital-twin and agent-based capabilities. A digital-twin simulation module, based on a classical Newtonian dissociation model, will predict experimental outcomes from user input and/or outputs of the Data Analysis & Featureization module. This module will guide the clustering process, and together they will form an interactive machine-learning loop that supports the validation and interpretation of measured results while identifying relevant features and underlying physical processes.

In particular, the agent will be employed to isolate and characterize dissociation dynamics in the concerted fragmentation pathway of water that remained inaccessible using conventional, by-hand experimental analysis. Only detailed *ab initio* theoretical investigations revealed that three of the eight water dication states exhibit non-standard dissociation behavior associated with a breakdown of the axial-recoil-approximation. This behavior arises from a so-called “slingshot mechanism,” which inverts the kinematics of the reaction products [43]. Notably, this effect could not be identified through tradi-

tional manual analysis of the experimental data alone.

Preliminary investigations using SCULPT indicate that the three clusters corresponding to these dication states exhibit varying degrees of overlap or separation in UMAP representations, depending on the choice of correlated features. The agent will be tasked with autonomously analyzing these UMAP patterns in conjunction with the digital-twin to identify shared characteristics among the clusters in feature space. This approach provides a data-driven pathway to detect events that violate the axial-recoil-approximation without relying on *ab initio* calculations, and it enables the systematic identification of non-axial fragmentation dynamics in complex molecular systems that are beyond the reach of such sophisticated theoretical methods.

In this regard, As such, the SCULPT platform exemplifies a transformative approach to analyzing high-dimensional tabulated data, with broad implications across both fundamental science and industry. As an outlook, We believe that this work lays the foundation for AI-driven discovery in multi-particle quantum dynamics, emphasizing real-time, physics-informed analysis of complex experimental datasets. In the long term, its modular, web-based architecture and adaptive clustering tools will not only accelerate insight in COLTRIMS experiments but also potentially offer scalable solutions for applications and industries where uncovering rare events and nonlinear correlations in large tabulated multi-parameter datasets are critical, such as quantum information science,[44] pharmaceuticals,[45–47] medicine,[48] aerospace. [49, 50] SCULPT can bridge experimental physics with data-centric innovation, positioning itself as a model for cross-sector impact. We have made SCULPT available on GitHub to encourage community contributions and to ensure that the software can evolve with the field’s needs.

ACKNOWLEDGMENTS

This work was supported by the Laboratory Directed Research and Development (LDRD) program at Lawrence Berkeley National Laboratory. We acknowledge the computational resources provided by the Advanced Light Source and the National Energy Research Scientific Computing Center (NERSC), which both are a DOE Office of Science User Facility under contract no. DE-AC02-05CH11231. In particular we acknowledge NERSC award ERCAP-0031498. We also thank the ALS Photon Science Computing Group for their valuable feedback.

DATA AVAILABILITY

The data that support the findings of this study are available from the corresponding author upon reasonable

request. Example datasets and tutorials are available at

<https://github.com/AMOS-experiment/CoInML/>.

- [1] R. Dörner, V. Mergel, O. Jagutzki, L. Spielberger, J. Ullrich, R. Moshhammer, and H. Schmidt-Böcking, Cold Target Recoil Ion Momentum Spectroscopy: a ‘momentum microscope’ to view atomic collision dynamics, *Physics Reports* **330**, 95 (2000).
- [2] J. Ullrich, R. Moshhammer, A. Dorn, R. Dörner, L. P. H. Schmidt, and H. Schmidt-Böcking, Recoil-ion and electron momentum spectroscopy: reaction-microscopes, *Reports on Progress in Physics* **66**, 1463 (2003).
- [3] T. Jahnke, A. Czasch, M. Schöffler, S. Schössler, M. Käs, J. Titze, K. Kreidi, R. E. Grisenti, A. Staudte, O. Jagutzki, L. P. H. Schmidt, S. K. Semenov, N. A. Cherepkov, H. Schmidt-Böcking, and R. Dörner, Photoelectron and ICD electron angular distributions from fixed-in-space neon dimers, *Journal of Physics B: Atomic, Molecular and Optical Physics* **40**, 2597 (2007).
- [4] M. S. Schöffler, J. Titze, N. Petridis, T. Jahnke, K. Cole, L. P. H. Schmidt, A. Czasch, D. Akoury, O. Jagutzki, J. B. Williams, N. A. Cherepkov, S. K. Semenov, C. W. McCurdy, T. N. Rescigno, C. L. Cocke, T. Osipov, S. Lee, M. H. Prior, A. Belkacem, A. L. Landers, H. Schmidt-Böcking, T. Weber, and R. Dörner, Ultrafast Probing of Core Hole Localization in N₂, *Science* **320**, 920 (2008).
- [5] T. Weber, A. O. Czasch, O. Jagutzki, A. K. Müller, V. Mergel, A. Kheifets, E. Rotenberg, G. Meigs, M. H. Prior, S. Daveau, A. Landers, C. L. Cocke, T. Osipov, R. Díez Muño, H. Schmidt-Böcking, and R. Dörner, Complete photo-fragmentation of the deuterium molecule, *Nature* **431**, 437 (2004).
- [6] K. Fehre, D. Trojanowskaja, J. Gatzke, M. Kunitski, F. Trinter, S. Zeller, L. P. H. Schmidt, J. Stohner, R. Berger, A. Czasch, O. Jagutzki, T. Jahnke, R. Dörner, and M. S. Schöffler, Absolute ion detection efficiencies of microchannel plates and funnel microchannel plates for multi-coincidence detection, *Review of Scientific Instruments* **89**, 045112 (2018).
- [7] P. V. Demekhin, S. D. Stoychev, A. I. Kuleff, and L. S. Cederbaum, Exploring Interatomic Coulombic Decay by Free Electron Lasers, *Physical Review Letters* **107**, 273002 (2011).
- [8] A. Czasch, M. Schöffler, M. Hattass, S. Schössler, T. Jahnke, T. Weber, A. Staudte, J. Titze, C. Wimmer, S. Kammer, M. Weckenbrock, S. Voss, R. E. Grisenti, O. Jagutzki, L. P. H. Schmidt, H. Schmidt-Böcking, R. Dörner, J. M. Rost, T. Schneider, C.-N. Liu, I. Bray, A. S. Kheifets, and K. Bartschat, Partial Photoionization Cross Sections and Angular Distributions for Double Excitation of Helium up to the N = 13 Threshold, *Physical Review Letters* **95**, 243003 (2005).
- [9] L. P. H. Schmidt, T. Jahnke, A. Czasch, M. Schöffler, H. Schmidt-Böcking, and R. Dörner, Spatial Imaging of the H₂ + Vibrational Wave Function at the Quantum Limit, *Physical Review Letters* **108**, 073202 (2012).
- [10] Y. LeCun, Y. Bengio, and G. Hinton, Deep learning, *Nature* **521**, 436 (2015).
- [11] H. Wang, T. Fu, Y. Du, W. Gao, K. Huang, Z. Liu, P. Chandak, S. Liu, P. Van Katwyk, A. Deac, A. Anandkumar, K. Bergen, C. P. Gomes, S. Ho, P. Kohli, J. Lasenby, J. Leskovec, T.-Y. Liu, A. Manrai, D. Marks, B. Ramsundar, L. Song, J. Sun, J. Tang, P. Veličković, M. Welling, L. Zhang, C. W. Coley, Y. Bengio, and M. Zitnik, Scientific discovery in the age of artificial intelligence, *Nature* **620**, 47 (2023).
- [12] Z. Chen, N. Andrejevic, N. C. Drucker, T. Nguyen, R. P. Xian, T. Smidt, Y. Wang, R. Ernstorfer, D. A. Tennant, M. Chan, and M. Li, Machine learning on neutron and x-ray scattering and spectroscopies, *Chemical Physics Reviews* **2**, 031301 (2021).
- [13] M. Hu, J. Fan, Y. Tong, Z. Sun, and H. Jiang, Deep learning for ultrafast X-ray scattering and imaging with intense X-ray FEL pulses, *Advanced Optical Technologies* **14**, 1546386 (2025).
- [14] H. Daoud, D. Sirohi, E. Mjeku, J. Feng, S. Oghbaey, and R. J. D. Miller, Novel applications of generative adversarial networks (GANs) in the analysis of ultrafast electron diffraction (UED) images, *The Journal of Chemical Physics* **159**, 044107 (2023).
- [15] C. Zhao, W. Yu, and L. Li, Visualization of small-angle X-ray scattering datasets and processing-structure mapping of isotactic polypropylene films by machine learning, *Materials & Design* **228**, 111828 (2023).
- [16] H. He, C. Liu, and H. Liu, Model Reconstruction from Small-Angle X-Ray Scattering Data Using Deep Learning Methods, *iScience* **23**, 100906 (2020).
- [17] T. Chavez, Z. Zhao, R. Jiang, W. Koepp, D. McReynolds, P. H. Zwart, D. B. Allan, E. H. Gann, N. Schwarz, D. Ushizima, E. S. Barnard, A. Mehta, S. Sankaranarayanan, and A. Hexemer, A machine-learning-driven data labeling pipeline for scientific analysis in *MLExchange*, *Journal of Applied Crystallography* **58**, 731 (2025).
- [18] W. M. Abdelmoula, B. G.-C. Lopez, E. C. Randall, T. Kapur, J. N. Sarkaria, F. M. White, J. N. Agar, W. M. Wells, and N. Y. R. Agar, Peak learning of mass spectrometry imaging data using artificial neural networks, *Nature Communications* **12**, 5544 (2021).
- [19] V. Bitto, P. Hönscheid, M. J. Besso, C. Sperling, I. Kurth, M. Baumann, and B. Brors, Enhancing mass spectrometry imaging accessibility using convolutional autoencoders for deriving hypoxia-associated peptides from tumors, *npj Systems Biology and Applications* **10**, 57 (2024).
- [20] L. McInnes, J. Healy, and J. Melville, UMAP: Uniform Manifold Approximation and Projection for Dimension Reduction (2020), arXiv:1802.03426 [stat].
- [21] E. Becht, L. McInnes, J. Healy, C.-A. Dutertre, I. W. H. Kwok, L. G. Ng, F. Ginhoux, and E. W. Newell, Dimensionality reduction for visualizing single-cell data using UMAP, *Nature Biotechnology* **37**, 38 (2019).
- [22] B. Szubert, J. E. Cole, C. Monaco, and I. Drozdov, Structure-preserving visualisation of high dimensional single-cell datasets, *Scientific Reports* **9**, 8914 (2019).
- [23] F. Trozzi, X. Wang, and P. Tao, UMAP as a Dimensionality Reduction Tool for Molecular Dynamics Simulations of Biomacromolecules: A Comparison Study, *The Journal of Physical Chemistry B* **125**, 5022 (2021).

- [24] D. Bourilkov, Machine and deep learning applications in particle physics, *International Journal of Modern Physics A* **34**, 1930019 (2019).
- [25] D. Guest, K. Cranmer, and D. Whiteson, Deep Learning and Its Application to LHC Physics, *Annual Review of Nuclear and Particle Science* **68**, 161 (2018).
- [26] J. Brehmer, S. Macaluso, D. Pappadopulo, and K. Cranmer, Hierarchical clustering in particle physics through reinforcement learning (2020), arXiv:2011.08191 [cs].
- [27] A. S. Venkatachalam, L. Greenman, J. Stallbaumer, A. Rudenko, D. Rolles, and H. V. S. Lam, enExploiting correlations in multi-coincidence Coulomb explosion patterns for differentiating molecular structures using machine learning, *Nature Communications* **16**, 11366 (2025).
- [28] X. Li, T. Jahnke, R. Boll, J. Han, M. Xu, M. Meyer, M. N. Piancastelli, D. Rolles, A. Rudenko, F. Trinter, T. J. A. Wolf, J. B. Thayer, J. P. Cryan, S. Ermon, and P. J. Ho, Generative Modeling Enables Molecular Structure Retrieval from Coulomb Explosion Imaging (2025), arXiv:2511.00179 [physics].
- [29] F.-A. Fortin, F.-M. De Rainville, M.-A. G. Gardner, M. Parizeau, and C. Gagné, DEAP: evolutionary algorithms made easy, *J. Mach. Learn. Res.* **13**, 2171 (2012).
- [30] Plotly Technologies Inc., Collaborative Data Science (2015).
- [31] N. Cherrier, J.-P. Poli, M. Defurne, and F. Sabatié, Consistent Feature Construction with Constrained Genetic Programming for Experimental Physics, in *2019 IEEE Congress on Evolutionary Computation (CEC)* (2019) pp. 1650–1658.
- [32] W. La Cava and J. H. Moore, Learning feature spaces for regression with genetic programming, *Genetic Programming and Evolvable Machines* **21**, 433 (2020).
- [33] N. Makke and S. Chawla, Interpretable scientific discovery with symbolic regression: a review, *Artificial Intelligence Review* **57**, 2 (2024).
- [34] D. Reedy, J. B. Williams, B. Gaire, A. Gatton, M. Weller, A. Menssen, T. Bauer, K. Henrichs, P. Burzynski, B. Berry, Z. L. Streeter, J. Sartor, I. Ben-Itzhak, T. Jahnke, R. Dörner, T. Weber, and A. L. Landers, Dissociation dynamics of the water dication following one-photon double ionization. II. Experiment, *Physical Review A* **98**, 053430 (2018).
- [35] P. J. Rousseeuw, Silhouettes: A graphical aid to the interpretation and validation of cluster analysis, *Journal of Computational and Applied Mathematics* **20**, 53 (1987).
- [36] B. Hopkins and J. G. Skellam, A New Method for determining the Type of Distribution of Plant Individuals, *Annals of Botany* **18**, 213 (1954).
- [37] R. G. Lawson and P. C. Jurs, New index for clustering tendency and its application to chemical problems, *Journal of Chemical Information and Computer Sciences* **30**, 36 (1990).
- [38] L. Hubert and P. Arabie, Comparing partitions, *Journal of Classification* **2**, 193 (1985).
- [39] T. Calinski and J. Harabasz, A dendrite method for cluster analysis, *Communications in Statistics - Theory and Methods* **3**, 1 (1974).
- [40] D. L. Davies and D. W. Bouldin, A Cluster Separation Measure, *IEEE Transactions on Pattern Analysis and Machine Intelligence* **PAMI-1**, 224 (1979).
- [41] O. Arbelaiz, I. Gurrutxaga, J. Muguerza, J. M. Pérez, and I. Perona, An extensive comparative study of cluster validity indices, *Pattern Recognition* **46**, 243 (2013).
- [42] M. Ester, H.-P. Kriegel, J. Sander, and X. Xu, A density-based algorithm for discovering clusters in large spatial databases with noise, in *Proceedings of the Second International Conference on Knowledge Discovery and Data Mining*, KDD'96 (AAAI Press, Portland, Oregon, 1996) pp. 226–231.
- [43] Z. L. Streeter, F. Yip, R. L. Lucchese, B. Gervais, T. N. Rescigno, and C. W. Mc McCurdy, Dissociation dynamics of the water dication following one-photon double ionization. I. Theory, *Physical Review A* **98**, 10.1103/PhysRevA.98.053429 (2018).
- [44] M. Sennary, J. Rivera-Dean, M. ElKabbash, V. Pervak, M. Lewenstein, and M. T. Hassan, Attosecond quantum uncertainty dynamics and ultrafast squeezed light for quantum communication, *Light: Science & Applications* **14**, 350 (2025).
- [45] J. A. Berlin, S. C. Glasser, and S. S. Ellenberg, Adverse Event Detection in Drug Development: Recommendations and Obligations Beyond Phase 3, *American Journal of Public Health* **98**, 1366 (2008).
- [46] C. Feng, L. Li, and A. Sadehpoor, A comparison of residual diagnosis tools for diagnosing regression models for count data, *BMC Medical Research Methodology* **20**, 175 (2020).
- [47] A. V. Sadybekov and V. Katritch, Computational approaches streamlining drug discovery, *Nature* **616**, 673 (2023).
- [48] Y. Huang, J. Li, M. Li, and R. R. Aparasu, Application of machine learning in predicting survival outcomes involving real-world data: a scoping review, *BMC Medical Research Methodology* **23**, 268 (2023).
- [49] A. Srivastava, R. Akella, V. Diev, S. Kumaresan, D. McIntosh, E. Pontikakis, Zuobing Xu, and Yi Zhang, Enabling the Discovery of Recurring Anomalies in Aerospace Problem Reports using High-Dimensional Clustering Techniques, in *2006 IEEE Aerospace Conference* (IEEE, Big Sky, MT, USA, 2006) pp. 1–17.
- [50] V. M. Janakiraman and D. Nielsen, Anomaly detection in aviation data using extreme learning machines, in *2016 International Joint Conference on Neural Networks (IJCNN)* (IEEE, Vancouver, BC, Canada, 2016) pp. 1993–2000.