Independent Component Analysis, Symmetry Reduction and Particle Physics

Kingsley J.A. Cox

Stony Brook, NY, USA

kcox@syndar.org

Abstract

Independent Component Analysis (ICA) explains the emergence of oriented receptive fields in visual cortex by uncovering higher-order statistical structure invisible to Principal Component Analysis (PCA). Whereas PCA preserves rotational symmetry in whitened data (SO(n)), ICA reduces this symmetry, producing a lobed energy landscape whose axes correspond to statistically independent sources—an analogue of spontaneous symmetry breaking. This paper interprets ICA as a progressive symmetry reduction process that, when applied hierarchically, parallels cortical processing from edges to objects and successive symmetry breakings in physics. We propose extending this framework to particle accelerator data, treating it as a high-dimensional mixture of latent physical sources. By applying recursive ICA, one could infer hidden subgroup structures within experimental data, tracing partial paths in a subgroup lattice to constrain possible parent gauge groups. This symmetry-based approach may provide a data-driven route toward uncovering deeper physical symmetries beyond the Standard Model.

1. Introduction

The success of Independent Component Analysis (ICA) in explaining the emergence of oriented, edge-like receptive fields in the visual cortex has made it one of the central models of sensory coding (Bell & Sejnowski, 1997; Hyvärinen & Hoyer, 2001; Simoncelli & Olshausen, 2001). When small patches of natural images are analyzed statistically, Principal Component Analysis (PCA) identifies directions of maximal variance—the second-order correlations—but fails to yield localized, oriented filters. ICA, in contrast, discovers oriented bar-like basis functions closely resembling simple-cell receptive fields in the primary visual cortex (V1). This difference arises because ICA exploits higher-order statistical structure in natural scenes that PCA cannot detect. The resulting ICA filters correspond to preferred orientations—directions in the space of image patches along which higher-order statistics are anisotropic. From a mathematical viewpoint, this can be interpreted as a symmetry reduction: the whitened PCA representation is isotropic, invariant under rotations, while ICA changes this isotropic sphere into a new 'lobed' shape whose axes define privileged directions.

Understanding ICA as a mechanism for symmetry reduction not only provides a new viewpoint in its role in cortical computation, but also suggests analogies to spontaneous symmetry breaking in physics and to the hierarchical emergence of structure in deep neural networks (Barlow, 1961; Bengio et al., 2013). The main idea presented in this paper is that symmetry reduction could be applied to particle accelerator data to uncover partial paths in a subgroup lattice, which may shed light on the fundamental group of how the universe operates. Any further information about this group, which is currently SU(3)×SU(2)×U(1) (the cornerstone of the Standard Model of physics), would be useful.

2. PCA, Natural Image Statistics, and the Isotropic Sphere

Natural image patches are not random pixel arrays but exhibit characteristic spatial correlations. If we take an image and extract small patches (e.g., 16 × 16 pixels), each patch can be viewed as a point in a high-dimensional space whose coordinates are the pixel intensities. Across a large ensemble of patches, the mean is typically removed, and PCA then identifies directions of maximal variance in this space (Jolliffe & Cadima, 2016).

PCA relies solely on second-order statistics—the covariance matrix ($C = E[xx^T]$). When the image patches are whitened (so that ($C = I$)), the data are represented as points on the surface of a hypersphere. In this space, there is no preferred orientation: all directions are statistically equivalent. In group-theoretic language, the system is invariant under the full rotation group SO(n). The PCA sphere thus represents rotational symmetry in the second-order structure of natural image patches (Simoncelli & Olshausen, 2001). However, visual experience tells us that the world is not rotationally symmetric—edges and contours occur more frequently along certain orientations, and local image structures are often elongated rather than isotropic. Yet these anisotropies do not appear at the level of pairwise correlations; they reside in higher-order dependencies that PCA cannot capture.

3. ICA and the Emergence of Oriented Bars

ICA extends beyond second-order correlations by maximizing the statistical independence of components. In its simplest linear form, we model the observed image patches (x) as mixtures of latent sources (s):

($x = As$)

where (A) is an unknown mixing matrix. The goal is to find an unmixing matrix ($W = A^{-1}$) such that the components ($y = Wx$) are as statistically independent as possible (Hyvärinen & Hoyer, 2001). While PCA decorrelates the data (removing second-order dependencies), ICA seeks to eliminate higher-order dependencies as well. The measure of independence can be expressed via non-Gaussianity—typically through kurtosis, negentropy, or contrast functions sensitive to higher-order cumulants (Hyvärinen et al., 2018). When Bell & Sejnowski (1997) applied ICA to natural image patches, the resulting filters were localized, oriented, and band-pass—closely matching the receptive fields of V1 simple cells. The algorithm had effectively discovered the statistical primitives of natural scenes: oriented edges or bars.

4. Higher-Order Statistics and Lobes

To visualize why ICA finds oriented bars while PCA does not, consider again the whitened patch space. After PCA, the data lie on a sphere - no preferred directions. When ICA maximizes non-

Gaussianity, it typically creates a lobed higher-order contrast landscape on the whitened sphere: the contrast function has multiple equivalent peaks (with sources having identical HOCS) aligned with source axes, producing a multi-lobed pattern of preferred directions corresponding to statistically independent sources. These axes are the independent components, analogous to orientations of edges in image space. The continuous rotational symmetry of the PCA sphere (SO(n)) is thus reduced to a finite subgroup consisting of discrete permutations, sign changes and rotations of the independent components (the symmetry operations of the lobed landscape).

4.1 Example: Symmetry Breaking with Three Identical Laplacian Sources

A particularly clear illustration of symmetry reduction in ICA arises when three sources are independent but statistically identical with respect to second order correlations - for example, three Laplacian (super-Gaussian) sources with equal kurtosis (s1, s2 and s3). In whitened coordinates, the data initially possess full rotational symmetry SO(3): any rotation in the three-dimensional feature space yields an equivalent mixture, since all three sources share the same marginal distribution. However, once the kurtosis-based ICA contrast function is applied, the potential landscape is no longer perfectly isotropic. Instead, it develops a six-lobed structure, with extrema aligned along the coordinate axes corresponding to the directions ±s1, ±s2 and ±s3 (figure 1). The continuous rotational symmetry SO(3) is thus reduced to the discrete dihedral group O_h, the symmetry group of a cube, which preserves the pattern of six equivalent lobes under 90° rotations and reflections. Each lobe represents an equivalent fixed point of the ICA dynamics, corresponding to one of the sign-permuted versions of the three sources. In this way, identical sources cause the ICA energy landscape to "crystallize" from a circularly symmetric potential into a cube-symmetric one—an analogue of spontaneous symmetry breaking in physical systems, where an initially isotropic state condenses into a discrete set of equivalent minima.

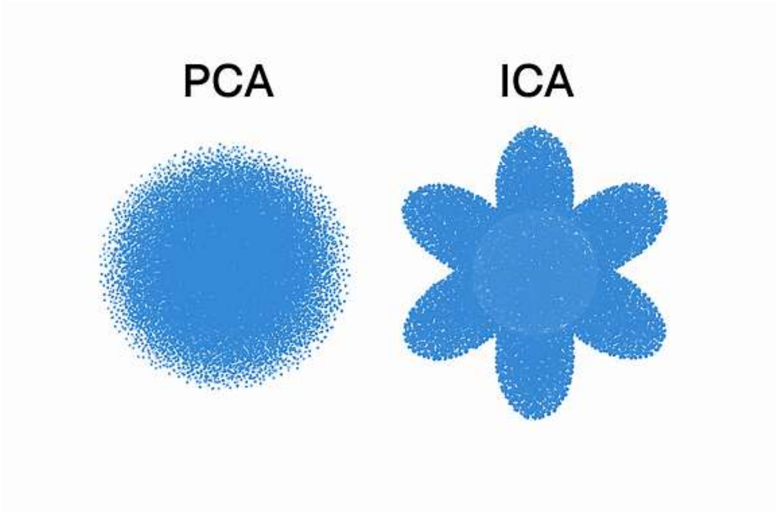

Figure 1 PCA and ICA symmetries

5. Symmetry Reduction and the Number/Type of Sources

Each stage of ICA can be viewed as a progressive symmetry reduction in feature space. Initially, the whitened data enjoy full SO(n) rotational invariance. The discovery of n independent components reduces this continuous symmetry to a discrete group of size 2^nxn!, corresponding to the possible sign flips and permutations of the sources. However, when subsets of sources share identical higher-order statistics, the symmetry is not fully broken: residual discrete symmetries persist within those subspaces. For example, if two sources are statistically indistinguishable, the latent space admits a D_4 symmetry — the dihedral group of order 8 — reflecting the ambiguity under 90° rotations and reflections. More generally, the final symmetry group includes both the global sign-permutation group and finite subgroups (such as D_k or S_k) acting within degenerate source clusters.

When a hierarchical or recursive ICA (Shan et al., 2006) is applied, each layer operates on outputs of the previous one, finding new independent directions in the transformed space. As the features become more structured—bars giving rise to corners, junctions, and object parts—the symmetry group contracts further. Each reduction corresponds to a restriction of allowable transformations: from continuous rotations to discrete permutations, and ultimately to the trivial group where each feature (object) is unique and cannot be transformed into another by symmetry.

In cortical terms, this hierarchy may correspond to the progression from V1 simple cells (oriented edges) to V2 cells (corners and textures) to inferotemporal neurons coding for specific objects (DiCarlo et al., 2012).

6. Group-Theoretic Interpretation

Mathematically, PCA yields a representation invariant under SO(n) rotations of the whitened data. ICA introduces a contrast function (J(W)) whose extrema correspond to directions of maximal non-Gaussianity. Each subgroup represents a particular symmetry reduction:

- PCA sphere → SO(n) symmetry.
- ICA lobed landscape → G = subgroup of SO(n).
- Hierarchical ICA → nested subgroups, smaller at each stage.

This sequence forms a **subgroup lattice** (Schmidt, 1994) (see section 11 below), analogous to the way physical symmetry groups break down under constraints or phase transitions.

7. Hierarchical ICA and Cortical Processing

Biological vision appears to implement such symmetry reductions hierarchically. In early visual cortex, neurons respond to localized oriented edges (Vinje & Gallant, 2000; Willmore et al., 2011). In higher visual areas, receptive fields integrate these components into more complex shapes (Hyvärinen & Hoyer, 2001).

At each stage, the network performs a new ICA on the outputs of the previous one, discovering higher-order structure in the activations rather than the raw pixels. The “basis vectors” at each level form a new coordinate frame in feature space, whose axes correspond to increasingly complex combinations of earlier features. Each stage reduces the symmetry of the representational manifold—akin to the way successive spontaneous symmetry breakings in physics lead from unified to specialized interaction laws (Croon et al., 2019).

8. Nonlinearity

If each layer of ICA were purely linear, successive transformations could be collapsed into a single linear transform. Thus, nonlinearity between layers is essential for generating genuinely new representations (Hyvärinen et al., 2018). In the brain, this role is played by neuronal nonlinearities - thresholding, rectification, and divisive normalization (Angelucci & Bressloff, 2017). These nonlinearities distort the distribution of responses, reintroducing higher-order structure for the next layer to exploit. Consequently, hierarchical ICA naturally mirrors cortical computation: a nonlinear encoding stage followed by a linear decorrelation (whitening) and independence-seeking stage. Each level's outputs form the effective "sources" for the next, producing a cascade of progressively reduced symmetries.

9. From Cortical ICA to Physics: Discovering Symmetry Groups

If ICA in the cortex amounts to symmetry reduction, then the same principle might reveal latent symmetries in other high-dimensional systems—such as fundamental physical data. Particle interactions recorded in accelerators produce vast multivariate datasets. Suppose these data are locally linear mixtures of a few underlying sources—latent fields or symmetries. A cortically inspired ICA network could attempt to recover those sources. The analogy runs deep. In particle physics, the Standard Model is structured around gauge symmetries—SU(3)×SU(2)×U(1)—that unify interactions (Kibble, 2014). Each symmetry-breaking stage (electroweak unification, confinement) corresponds to a reduction in symmetry similar to hierarchical ICA stages. If we feed raw experimental data into a network performing ICA or its nonlinear extensions, the emergent basis directions might correspond to hidden group generators or conserved quantities (Gabel et al., 2023; Laird et al., 2025).

In this framework, the initial PCA sphere represents a fully symmetric field—analogous to an unbroken gauge group. ICA's discovery of preferred directions represents the breaking of this symmetry, isolating physically meaningful degrees of freedom. Recursive ICA could, in principle, reveal further subgroup structures—mirroring successive physical symmetry breakings from unified to observable particles.

10. Experimental Possibilities

A "particle-ICA" experiment would begin by segmenting raw detector data into spatiotemporal "patches," analogous to visual image patches, small enough that local linear mixtures are plausible. After whitening, one would perform ICA to identify statistically independent directions—interpretable as independent dynamic modes. If the data are truly isotropic, the ICA landscape will be spherical and degenerate; any stable fixed points found would correspond to spontaneous alignment directions chosen by the algorithm.

If structure exists—say, anisotropies in higher-order correlations reflecting specific particle interactions—then the learned axes might correspond to physical symmetries or conserved quantities. By repeating the procedure hierarchically, one might recover subgroup hierarchies resembling those predicted by grand-unified models (Croon et al., 2019).

This approach connects representational learning in neuroscience with the algebraic structure of physics: both can be seen as processes that reduce symmetry to discover structure.

Particle accelerators, such as the Large Hadron Collider, produce high-energy collisions whose outcomes are recorded by complex detectors. These detectors do not output simple numbers but vast structured data describing where, when, and how much energy particles deposit as they pass through different layers. Each collision event generates millions of readings — positions, times, and energies — that are processed into higher-level objects like particle tracks, calorimeter clusters, and energy sums.

The resulting datasets can be represented as tensors, images, point clouds, or graphs, making them well suited to neural network analysis. Modern particle physics increasingly relies on machine learning for tasks such as particle identification, jet classification, event tagging, and anomaly detection. Convolutional, graph, and transformer networks are particularly effective at recognizing spatial and relational patterns in these data.

By converting raw detector hits into structured inputs — for example, 3D energy maps or sets of particle momenta — neural networks can learn to reconstruct and interpret events automatically, replacing or augmenting traditional rule-based methods. In this way, accelerator data provide some of the most complex and physics-rich inputs used in machine learning, bridging experimental particle physics and artificial intelligence.

The physics community **is** using neural networks to uncover symmetry groups in accelerator-type data. For example, SymmetryGAN (Desai et al. 2022) applied a GAN-based method to simulated LHC dijet events and was able to infer the underlying symmetry group of the dataset rather than assuming it a priori.

11. Future Prospects: Subgroup Lattices as Clues to Underlying Symmetries

Finally, this framework suggests a broader application beyond cortical computation. The same logic of progressive symmetry reduction that characterizes hierarchical ICA could serve as a methodological guide for discovering hidden symmetries in physics. If independent - analyses perhaps using cortically inspired ICA applied to diverse experimental datasets - reveal partial or overlapping paths of symmetry reduction, these can be interpreted as trajectories through the subgroup lattice of an as-yet-unknown group **G**. Even fragmentary subgroup chains encode valuable structural information: overlaps between them constrain the architecture of possible parent groups. Thus, convergent evidence from multiple ICA-based decompositions could allow physicists to triangulate the identity of **G**, extending beyond the Standard Model by reconstructing the unique higher-symmetry group whose subgroup lattice is consistent with all observed reduction paths (Schmidt, 1994; Croon et al., 2019; Armenta & Jodoin, 2020). In this view, the brain's strategy of learning through symmetry reduction becomes not only a model of perception but also a possible scalable epistemic instrument for revealing the deep group structure of physical law.

12. Conclusion

ICA's success in modelling cortical receptive fields derives from its ability to exploit higher-order statistics that break the isotropy of natural image patches. Conceptually, ICA transforms the PCA sphere—uniform and symmetric—into a lobed structure whose axes represent meaningful directions in feature space. This geometric deformation corresponds to a reduction in symmetry from the continuous group SO(n) to a discrete subgroup (in the case of 2 sources with equal kurtoses we would get the group D_4, the symmetry group of the square).
Each subsequent ICA layer introduces further reductions, creating new bases aligned with increasingly complex features. The progression from edges to objects mirrors a sequence of group-theoretic contractions, culminating in the trivial symmetry of unique objects. The same principle can, in theory, be applied to particle-physics data: starting from isotropic fields and revealing hidden symmetries through hierarchical symmetry breaking.
To rephrase, we apply this idea to high-energy physics data, with the potential to *infer a gauge group different from the Standard Model*. We propose a bottom-up approach: find small subgroups in event correlations, and then see if they "fit together", using partial paths in a subgroup lattice, into a larger symmetry G.

References

Angelucci A., & Bressloff P. C. (2017). Circuits and mechanisms for surround modulation in visual cortex. *Annual Review of Neuroscience,* 40, 425–451. https://doi.org/10.1146/annurev-neuro-072116-031418

Armenta M. A., & Jodoin P. (2020). *The Representation Theory of Neural Networks.* arXiv:2007.12213.

Barlow H. (1961). Possible Principles Underlying the Transformations of Sensory Messages. *Sensory Communication,* 1, 217–234. https://doi.org/10.7551/mitpress/9780262518420.003.0013

Barlow H. B. (2001). Redundancy reduction revisited. *Network: Computation in Neural Systems,* 12(3), 241–253.

Bell A. J., & Sejnowski T. J. (1997). The independent components of natural scenes are edge filters. *Vision Research,* 37(23), 3327–3338.

Bengio Y., Courville A., & Vincent P. (2013). Representation learning: A review and new perspectives. *IEEE TPAMI.*

Chalk M., Marre O., & Tkačik G. (2018). Toward a unified theory of efficient, predictive, and sparse coding. *PNAS,* 115(1), 186–191.

Croon D., Gonzalo T. E., Graf L., Košnik N., & White G. (2019). GUT Physics in the Era of the LHC. *Frontiers in Physics,* 7, 76.

Desai K., Nachman B., & Thaler J. (2022). SymmetryGAN: Symmetry discovery with deep learning. *Physical Review D*, 105(9), 096031.

DiCarlo J. J., Zoccolan D., & Rust N. C. (2012). How does the brain solve visual object recognition? *Neuron,* 73(3), 415–434.

Gabel A., Klein V., Valperga R., Webster J. S. W., Quax R., & Gavves E. (2023). Learning Lie Group Symmetry Transformations with Neural Networks. arXiv:2307.01583.

Hyvärinen A., & Hoyer P. O. (2001). A two-layer sparse coding model learns simple and complex cell receptive fields. *Vision Research,* 41(18), 2413–2423.

Hyvärinen A., Sasaki H., & Turner R. E. (2018). Nonlinear ICA using auxiliary variables and generalized contrastive learning. arXiv:1805.08651.

Jolliffe I. T., & Cadima J. (2016). Principal component analysis: A review and recent developments. *Philosophical Transactions of the Royal Society A,* 374(2065).

Kibble T. (2014). *The Standard Model of Particle Physics. European Review,* 23.

Laird L., Hsu C., Bapat A., & Walters R. (2025). MatrixNet: Learning over symmetry groups using learned group representations. arXiv:2501.09571