



Causal Deep Learning

M. Alex O. Vasilescu  

IPAM, University of California, Los Angeles CA, USA
Tensor Vision, Los Angeles CA, USA

Abstract. We derive a set of causal deep neural networks whose architectures are a consequence of tensor (multilinear) factor analysis, a framework that facilitates causal inference. Forward causal questions are addressed with a neural network architecture composed of causal capsules and a tensor transformer. Causal capsules compute a set of invariant causal factor representations, whose interactions are governed by a tensor transformation. Inverse causal questions are addressed with a neural network that implements the multilinear projection algorithm. The architecture reverses the order of the operations of a forward neural network and estimates the causes of effects. As an alternative to aggressive bottleneck dimension reduction or regularized regression that may camouflage an inherently underdetermined inverse problem, we prescribe modeling different aspects of the mechanism of data formation with piecewise tensor models whose multilinear projections produce multiple candidate solutions. Our forward and inverse questions may be addressed with shallow architectures, but for computationally scalable solutions, we derive a set of deep neural networks by taking advantage of block algebra. An interleaved kernel hierarchy results in a doubly non-linear tensor factor models. The causal neural networks that are a consequence of tensor factor analysis are data agnostic, but are illustrated with facial images. Sequential, parallel and asynchronous parallel computation strategies are described.

Keywords: factor analysis · explanatory · confirmatory · latent variables · causality · tensor algebra · deep learning · generative · discriminant

1 Introduction

Neural networks are being employed increasingly in high-stakes application areas, such as face recognition [70,33,69,14,93], and medical diagnosis [40,51,73]. Developing neural networks that offer causal explanations for correct results or failures is crucial in establishing trustworthy artificial intelligence.

Causal explanations specify the causes, the mechanism, and the conditions for replicating an observed effect [47,24,92]. Quantitatively, causality is the direct relationship between two events, A and B, where “ $A \xrightarrow{\text{causes}} B$ ” means “the effect of A is B”, a measurable and experimentally repeatable phenomena. Once verified with either experimental or observational studies, the statement “the effect of A is B” stays true regardless of new discoveries and changes in knowledge [31]. Causal explanations are a factual understanding, whereas interpretations – a concept also advocated for the development of trustworthy artificial intelligence [48] – are an understanding relative to a specific reference frame, such as a particular point of view or knowledge base. As new knowledge emerges, prior interpretations may be deemed to be inaccurate or invalid, which can undermine their reliability and usefulness in the development of trustworthy artificial intelligence. Interpretations are open to reinterpretation. Therefore, explanations should not be confused with interpretations or speculations masquerading as explanations [49,48].

* To appear doi.org/10.1007/978-3-031-78189-6_27.

The validity and robustness of causal explanations depend on causal model specifications in conjunction with the experimental designs used for acquiring training data [63]. Generative artificial intelligence research that focuses on unsupervised deep neural networks is not well suited for drawing causal conclusions. (Determining the causal factors from unsupervised training data is an ill-posed inverse problem since different combinations of the same causal factors can result in the same outcome.) Prior unsupervised deep neural network research [21,7,54,71] briefly addressed the connection to tensor (multilinear) factor analysis. Tensor (multilinear) factor analysis is a supervised learning approach that models the causal mechanisms that generate data [82,85], estimates the effects of causes [84,89], as well as the causes of effects, given an estimated forward model and constraints on the solution set [86,78]. Data and its representations – such as those derived from principal or independent component analysis etc. – are factorized into a set of invariant, discriminant, and disentangled causal representations [77][4.15][82]. These causal representations maximize the ratio of inter- to intra-class scatter [82]. Thus, a causal tensor model is not only generative, but discriminative. In fact, the most successful deep neural network approaches are also supervised with generative and/or discriminative properties [43,70,53]. However, their goal has been prediction rather than causal inference.

Consequently, we derive a set of causal deep neural networks that are a consequence of causal tensor factor analysis, Figs. 1-5. Tensor factor analysis is a transparent framework for both forward [84,82] and inverse causal inference [86,78].¹

Forward causal inference is a hypothesis-driven process, as opposed to a data-driven process, that models the mechanism of data formation and estimates the effects of interventions or counterfactuals [58,36,68,80]. Inverse causal “inference” estimates the causes of effects given an estimated forward model and constraints on the solution set [78,26]. By comparison, conventional statistics and machine learning model data distributions, predict a variable co-observed with another, or perform time series forecasting.

1.1 Causal Inference Versus Regression

Neural networks and tensor factorization methods may perform causal inference, or simply perform regression from which no causal conclusions are drawn. For causal inference, model specifications (*i.e.*, problem definition) and experimental design for acquiring training data trump algorithmic or neural network design and analysis, Fig. 2.

Causal tensor factor analysis was employed in the analysis and recognition of facial identities [82,80] with sparseness constraints [27,61], as well as in analysis and recognition of facial expressions [34], human motion signatures [17,75], and 3D sound [28]. It was employed in the rendering of arbitrary scenes, views and illuminations [84], and in the transfer of facial expressions [89] etc. Tensor factor analysis was also employed in psychometrics [74,29,13,8,44], econometrics [39,52], chemometrics [10], and signal processing [18]. Simple tensor regression and decompositions which do not draw causal conclusions, leveraged row, column and fiber redundancies to estimate missing data [15] and to perform rank reduction [95,90,11,38,32,6]. Recently, tensor dimensionality reduction and contractions have been employed in machine learning to reduce neural network parameters. Network parameters are organized into “data tensors”, and dimensionally reduced [46,55,42,41,57] or efficiently contracted [22].

¹ TensorFaces is a gentle introduction to causal tensor factor analysis [82,85].

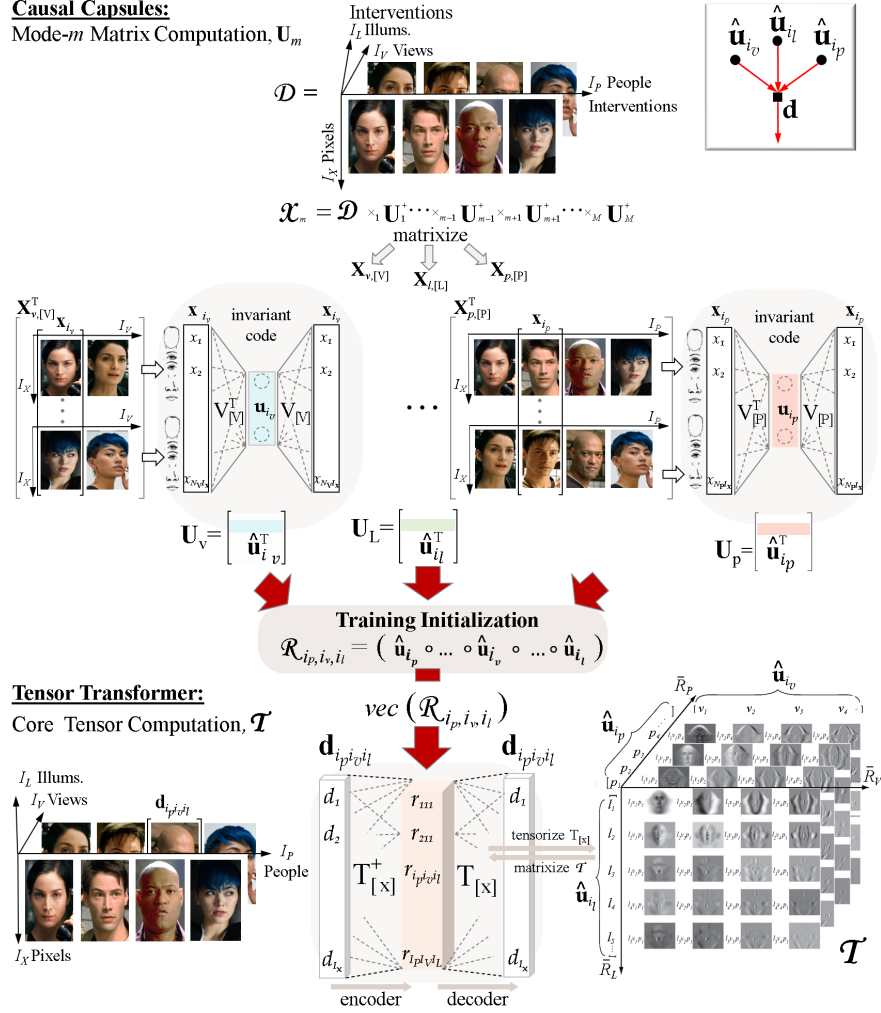


Fig. 1: A causal neural network that is an instantiation of the M -mode SVD [82,85], Alg. 1. A causal neural network is composed of a set of causal capsules and a tensor transformer. Given a set of vectorized and centered training observations organized into a “data tensor” \mathcal{D} , causal capsules learn a set of M subspaces $\{\mathbf{U}_m | 1 \leq m \leq M\}$ that span the M causal factor representations. The tensor transformer estimates the extended core \mathcal{T} , which governs the interaction between causal factors. A forward causal model employs the estimated tensor decoder decoupled from the encoder to generate a new data vector. A new data vector \mathbf{d}_{new} is the effect of new interventions, $\text{vec } \mathcal{R} = (\otimes_{m=1}^M \hat{\mathbf{u}}_{m_{\text{new}}})$, where \otimes is the Kronecker product. For scalable computations, each shallow autoencoder-decoder is replaced by a mathematically equivalent deep network, derived with the aid of block algebra, Fig. 3. In the upper-right corner, the causal model is summarized with graphical tensor notation [60], which may also be interpreted as a structural causal model DAG [58]. For display purposes only, the images were displayed in grid form and without being mean centered.

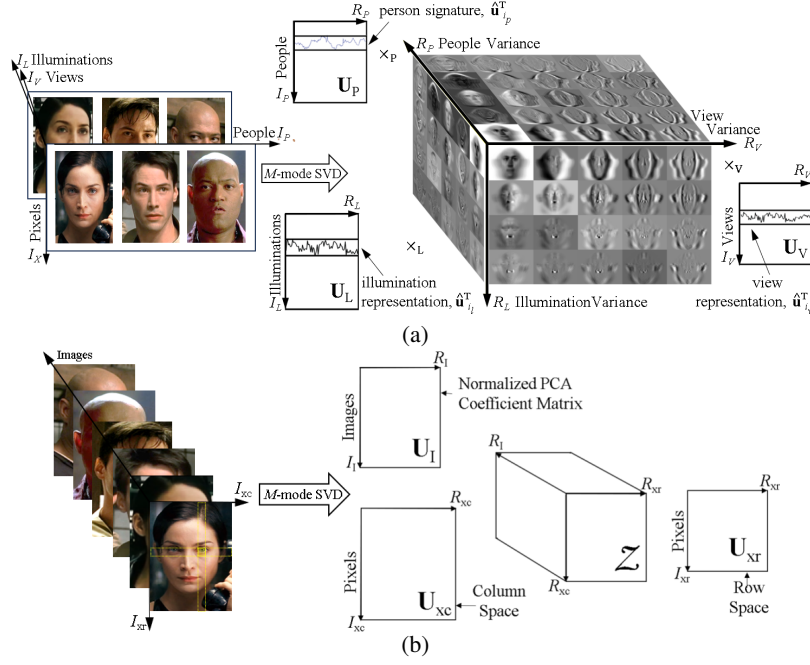


Fig. 2: Same data, same algorithm, but two different model specifications (problem setups / definitions). (a) Causal Inference: The M -mode SVD (Alg. 1) factorizes a “data tensor” of vectorized observations into multiple subspace that span the causal factor representations. (b) Simple regression: The M -mode SVD factorizes a “data tensor” composed of images as a 2-way array of column or row observations into a column and row space, as well as into normalized PCA coefficients. All images are vectorized except in Fig. 2b.

1.2 Causal Neural Networks

Causal neural networks are composed of causal capsules and a tensor transformer, Fig. 1. *Causal capsules* estimate the latent variables that represent the causal factors of data formation, and a *tensor transformer* governs their interaction. Causal capsules may be shallow autoencoder-decoder architectures that employ linear neurons and compute a set of invariant representations [67,65,64,1,56], as detailed in Supplemental A. The tensor transformer may be a *tensor autoencoder-decoder*, a shallow autoencoder-decoder whose code is the tensor product of the latent variables.

Causal deep neural networks are composed by stacking autoencoders-decoders. Each autoencoder-decoder in a shallow causal neural network is replaced by mathematically equivalent deep neural network architectures that are derived by taking advantage of block algebra. An interleaved hierarchy of kernel functions [66] serves as a pre-processor that warps the data manifold for optimal tensor factor analysis. A part-based deep neural network mirrors a part-based hierarchy of tensor factor models [80,79][77, Sec 4.4], Supplemental C.²

² There have been a number of related transformer architectures engineered and empirically tested with success [23,91,50].

Algorithm 1 M -mode SVD (parallel computation)[83,82]**Input** $\mathcal{D} \in \mathbb{C}^{I_0 \times \dots \times I_M}$, dimensions $R_0, R_1 \dots R_m \dots R_M$ 1. Initialize $\mathbf{U}_m := \mathbf{I}$, $0 \leq m \leq M$

2. Iterate until convergence

For $m := 0, \dots, M$,2.1 $\mathcal{X} := \mathcal{D} \times_0 \mathbf{U}_0^T \times \dots \times_{m-1} \mathbf{U}_{m-1}^T \times_{m+1} \mathbf{U}_{m+1}^T \dots \times \mathbf{U}_M^T$ 2.2 Set \mathbf{U}_m to the \tilde{R}_m leading left-singular vectors of the SVD of $\mathbf{X}_{[m]}$ or SVD of $[\mathbf{X}_{[m]} \mathbf{X}_{[m]}^T]^{a, b}$ 3. Set $\mathcal{Z} := \mathcal{D} \times_0 \mathbf{U}_0^T \dots \times_m \mathbf{U}_m^T \dots \times_M \mathbf{U}_M^T := \mathcal{X} \times \mathbf{U}_M^T$ ^c**Output** mode matrices $\mathbf{U}_0, \mathbf{U}_1, \dots, \mathbf{U}_M$ and core tensor \mathcal{Z} .

^a The computation of \mathbf{U}_m in the SVD $\mathbf{X}_{[m]} = \mathbf{U}_m \mathbf{\Sigma} \mathbf{V}_m^T$ can be performed efficiently, depending on which dimension of $\mathbf{X}_{[m]}$ is smaller, by decomposing either $\mathbf{X}_{[m]} \mathbf{X}_{[m]}^T = \mathbf{U}_m \mathbf{\Sigma}^2 \mathbf{U}_m^T$ (note that $\mathbf{V}_m^T = \mathbf{\Sigma}^+ \mathbf{U}_m^T \mathbf{X}_{[m]}$) or by decomposing $\mathbf{X}_{[m]}^T \mathbf{X}_{[m]} = \mathbf{V}_m \mathbf{\Sigma}^2 \mathbf{V}_m^T$ and then computing $\mathbf{U}_m = \mathbf{X}_{[m]} \mathbf{V}_m \mathbf{\Sigma}^+$.

^b For a neural network implementation, the SVD of $\mathbf{X}_{[m]}$ is replaced with an autoencoder that sequentially computes the orthonormal columns of $\mathbf{U}_m / \mathbf{V}_m$ by performing gradient descent with the learning parameter η or stochastic gradient descent [9,62]. In Fig. 1, the autoencoders compute the columns of \mathbf{V}_m , where $\mathbf{v}_{m,r}$ is the r column and it represents the weights of the r neuron. Matrix $\mathbf{V}_{m,r}$ contains the first r columns of \mathbf{V}_m .

$$\left. \begin{array}{l} \text{Autoencoder} \\ \left\{ \begin{array}{l} \text{For } r := 1 \dots R_m. \\ \text{Iterate until convergence} \\ \Delta \mathbf{v}_{m,r}(t+1) = \eta (\mathbf{X}_{[m]} - \mathbf{V}_{m,r-1}(t) \mathbf{V}_{m,r-1}^T(t) \mathbf{X}_{[m]}) \underbrace{\mathbf{X}_{[m]}^T \mathbf{v}_{m,r}(t)}_{\text{code}} \\ \hat{\mathbf{v}}_{m,r}(t+1) = \frac{(\mathbf{v}_{m,r}(t) + \Delta \mathbf{v}_{m,r}(t+1))}{\|\mathbf{v}_{m,r}(t) + \Delta \mathbf{v}_{m,r}(t+1)\|} \end{array} \right. \end{array} \right\}$$

^c The columns in $\mathbf{Z}_{[0]}$ may be computed by initializing the code of an autoencoder to $(\mathbf{U}_M \dots \otimes \mathbf{U}_m \dots \otimes \mathbf{U}_0)$, where \otimes is the Kronecker product. In Fig. 1, the columns of the extended core \mathcal{T} are computed by initializing the code of the autoencoder with $(\mathbf{U}_M \dots \otimes \mathbf{U}_m \dots \otimes \mathbf{U}_1)^T$ for batch training, and $(\hat{\mathbf{u}}_{i_M}^T \dots \otimes \hat{\mathbf{u}}_{i_m}^T \dots \otimes \hat{\mathbf{u}}_{i_1}^T)^T$ when training one observation, $\mathbf{d}_{i_1, \dots, i_M}$, at a time.

Inverse causal neural networks implement the multilinear projection algorithm to estimate the causes of effects [86,78]. A neural network that addresses an underdetermined inverse problem is characterized by a wide hidden layer. Dimensionality reduction removes noise and nuisance variables [30,72], and has the added benefit of reducing the widths of hidden layers. However, aggressive bottleneck dimensionality reduction may camouflage an inherently ill-posed problem. Alternatively or in addition to dimensionality reduction and regularized regression, we prescribe modeling different aspects of the data formation process with piecewise tensor (multilinear) models that return a set of candidate solutions [81]. Candidate solutions are gated to yield the most likely solution.

2 Forward Causal Question: “What if?”

Forward causal inference is a hypothesis-driven process that addresses the “what if” question. What if A is changed by one unit, what is the expected change in B? Causal hypotheses drive both the model specification and the experimental design for acquiring or generating training data.

2.1 Training Data

For modeling unit level effects of causes, the training data is generated by combinatorially varying each causal factor while holding the other factors fixed. The best causal evidence comes from randomized experimental studies. When randomized experiments for generating training data are unethical or infeasible, experimental studies may be approximated with carefully designed observational studies [63], such as natural experiments [2, 12, 37] or by employing the concept of transportability to transfer information from a source population where both experimental and observational studies are available to a target population where only observational studies can be conducted [59].³

2.2 Tensor Factor Analysis Model

Within the tensor mathematical framework, an $(M+1)$ -way array $\mathcal{D} \in \mathbb{C}^{I_0 \times I_1 \times \dots \times I_M}$ is a collection of vectorized and centered observations $\{\mathbf{d}_{i_1 \dots i_M} \in \mathbb{C}^{I_0}\}$ generated by M causal factors. Causal factor m , for $m = 1 \dots M$, takes one of I_m possible values, indexed by i_m , where $i_m = 1 \dots I_m$. An observation and an array of observations may be modeled using multilinear (tensor) principal component analysis,

$$\mathbf{d}_{i_1, \dots, i_M} = \mathcal{T} \times_1 \hat{\mathbf{u}}_{i_1}^T \cdots \times_M \hat{\mathbf{u}}_{i_M}^T + \mathbf{e}_{i_1, \dots, i_M}, \quad (1)$$

$$\mathcal{D} = \mathcal{T} \times_1 \mathbf{U}_1 \cdots \times_M \mathbf{U}_M + \mathcal{E}, \quad (2)$$

where the extended core tensor \mathcal{T} governs the interaction between the latent variables $\{\hat{\mathbf{u}}_{i_m} \in \mathbb{C}^{\tilde{I}_m} | 1 \leq m \leq M \text{ and } \tilde{I}_m \leq I_m\}$ that represent the causal factors of data formation, and $\hat{\mathbf{u}}_{i_m}^T$ is a row in $\mathbf{U}_m \in \mathbb{C}^{I_m \times \tilde{I}_m}$, Fig. 2a.⁴ The extended core tensor $\mathcal{T} = \mathcal{D} \times_1 \mathbf{U}_1^T \cdots \times_M \mathbf{U}_M^T = \mathcal{Z} \times_0 \mathbf{U}_0$ contains the vectors that span observation subspace.⁵ The error $\mathbf{e}_{i_1, \dots, i_M} \in \mathbb{C}^{I_0}$ is a vector in $\mathcal{E} \in \mathbb{C}^{I_0 \times I_1 \times \dots \times I_M}$ and $\mathbf{e}_{i_1, \dots, i_M} = \mathcal{Z} \times_0 \mathbf{E}_0 \times_1 \boldsymbol{\epsilon}_{i_1}^T \cdots \times_M \boldsymbol{\epsilon}_{i_M}^T$, where $\boldsymbol{\epsilon}_{i_m}$ drawn from a Gaussian distribution $\mathcal{N}(\mathbf{0}, \boldsymbol{\Sigma}_m)$ associated with each mode m ranging from 0 to M and $\mathbf{E}_0 = [\dots \boldsymbol{\epsilon}_{i_0} \dots \boldsymbol{\epsilon}_{I_0}]^T$.

Minimizing the cost function

³ Gebru *et al.*’s “Datasheets for datasets” may aid in the approximation of experimental studies [25].

⁴ A tensor \mathcal{T} is a multilinear mapping from a set of domain vector spaces to a range vector space, $\mathcal{T} : \{\mathbb{C}^{I_1} \times \dots \times \mathbb{C}^{I_M}\} \rightarrow \mathbb{C}^{I_0}$. In a causal tensor framework, the M domain spaces span the causal factor representations and the range vector space spans the observation space. An M -way array of vectorized observations, \mathcal{D} , is sometimes informally referred to as a “data tensor”, but it is not an actual tensor. It is preferable to vectorize an image and treat it as a single observation rather than as a collection of independent column/row observations [77, App. A]. For a basic review of tensor algebra, see Supplemental B.

⁵ In practice, the measurement mode \mathbf{U}_0 , *i.e.*, the PCA basis matrix, and the core tensor \mathcal{Z} are not computed.

$$L = \|\mathcal{D} - \mathcal{T} \times_1 \mathbf{U}_1 \dots \times_m \mathbf{U}_m \dots \times_M \mathbf{U}_M\| + \sum_{m=1}^M \lambda_m \|\mathbf{U}_m^T \times_1 \mathbf{U}_m - \mathbf{I}\| \quad (3)$$

is equivalent to maximum likelihood estimation [20] of the causal factor parameters, assuming the data was generated by the model with additive Gaussian noise. The cost function (3) is minimized by employing M alternating least squares optimizations [19, 44]

$$L_m = \|\mathcal{X}_m - \mathcal{T} \times_m \mathbf{U}_m\| + \lambda_m \|\mathbf{U}_m^T \times_1 \mathbf{U}_m - \mathbf{I}\|, \quad (4)$$

where

$$\mathcal{X}_m := \mathcal{D} \times_1 \dots \times_{m-1} \mathbf{U}_{m-1}^T \times_{m+1} \mathbf{U}_{m+1}^T \dots \times_M \mathbf{U}_M^T \quad \text{- parallel computation of all } \mathcal{X}_m \quad (5)$$

$$= \mathcal{X}_m(t-1) \times_n \mathbf{U}_n^T(t) \mathbf{U}_n(t-1), \quad \forall n \neq m, \quad \text{- asynchronous parallel computation} \quad (6)$$

$$= (\mathcal{X}_{m-1} \times_{m-1} \mathbf{U}_{m-1}^T) \times_m \mathbf{U}_m \quad \text{- sequential computation} \quad (7)$$

The M -mode SVD [82] (Alg. 1) minimizes the M alternating least squares (4) in closed form by employing M different SVDs. The approach is suitable for parallel (5), asynchronous (6), or sequential (7) computation. The extended core tensor \mathcal{T} is computed by multiplying the data tensor with the inverse mode matrices, $\mathcal{T} = \mathcal{D} \times_1 \mathbf{U}_1^T \dots \times_m \mathbf{U}_m^T \dots \times_M \mathbf{U}_M^T$, or more efficiently as $\mathcal{T} = \mathcal{X}_m \times \mathbf{U}_m^T$ where \mathcal{X}_m is the last updated one.

2.3 Kernel Tensor Factor Analysis Model

When data \mathcal{D} is a tensor combination of non-linear independent causal factors, kernel multilinear independent component analysis (K-MICA) [77, Ch 4.4] employs the “kernel trick” [66, 88] as a pre-processing step which makes the data suitable for multilinear independent component analysis [85] (Alg. 2),

$$\mathcal{D} = \mathcal{T} \times_1 \mathbf{C}_1 \dots \times_m \mathbf{C}_m \dots \times_M \mathbf{C}_M + \mathcal{E} \quad (8)$$

$$\mathbf{C}_m = \mathbf{U}_m \mathbf{W}_m^{-1} \quad (9)$$

where \mathbf{C}_m are the independent components, \mathbf{W}_m is a rotation matrix computed either with mutual information, negentropy or higher-order cumulants, and $\mathcal{E} = \mathcal{Z} \times_0 \mathbf{E}_0 \times_1 \mathbf{E}_1 \dots \times_M \mathbf{E}_M$ is the error. K-MPCA is a tensor generalization of the kernel PCA [66] and K-MICA is a tensor generalization of kernel ICA [3, 94].

To accomplish this analysis, recall that the computation of covariance matrix $\mathbf{D}_{[m]} \mathbf{D}_{[m]}^T$ involves inner products $\mathbf{d}_{i_1 \dots i_{m-1} j i_{m+1} \dots i_M}^T \mathbf{d}_{i_2 \dots i_{m-1} k i_{m+1} \dots i_M}$ between pairs of data points in the data tensor \mathcal{D} associated with causal factor m , for $m = 1, \dots, M$ (Step 2.2 in Algorithm 1). We replace the inner products with a generalized distance measure between images, $K(\mathbf{d}_{i_1 \dots i_{m-1} j i_{m+1} \dots i_M}, \mathbf{d}_{i_2 \dots i_{m-1} k i_{m+1} \dots i_M})$, where $K(\cdot, \cdot)$ is a suitable kernel function (Table 1) that corresponds to an inner product in some expanded feature space. This generalization naturally leads us to a *Kernel Multilinear PCA (K-MPCA) Algorithm*, where the covariance computation is replaced by

$$[\mathbf{D}_{[m]} \mathbf{D}_{[m]}^T]_{jk} := \sum_{i_1=1}^{I_1} \dots \sum_{i_{m-1}=1}^{I_{m-1}} \sum_{i_{m+1}=1}^{I_{m+1}} \dots \sum_{i_M=1}^{I_M} K(\mathbf{d}_{i_1 \dots i_{m-1} j i_{m+1} \dots i_M}, \mathbf{d}_{i_1 \dots i_{m-1} k i_{m+1} \dots i_M}). \quad (10)$$

Algorithm 2 Kernel Tensor Factor Analysis [77, Sec 4.4],[85]

Kernel Multilinear Independent Component Analysis (K-MICA) and
Kernel Principal Component Analysis (K-MPCA).

Input the data tensor $\mathcal{D} \in \mathbb{C}^{I_0 \times \dots \times I_M}$, where mode $m = 0$ is the measurement mode, and the desired ranks are $\tilde{R}_1, \dots, \tilde{R}_M$.

Initialize $\mathbf{C}_m = \mathbf{I}, \forall 0 \leq m \leq M$

Iterate until convergence.

1. For $m := 1, \dots, M$
 - (a) Set $\mathcal{X}_m := \mathcal{D} \times_1 \mathbf{C}_1^+ \dots \times_{m-1} \mathbf{C}_{m-1}^+ \times_{m+1} \mathbf{C}_{m+1}^+ \dots \times_M \mathbf{C}_M^+$.
 - (b) Compute the elements of the mode- m covariance matrix using kernel functions, Table 1, for $j, k := 1, \dots, I_m$:

$$[\mathbf{X}_{[m]} \mathbf{X}_{[m]}^T]_{jk} := \sum_{i_1=1}^{I_1} \dots \sum_{i_{m-1}=1}^{I_{m-1}} \sum_{i_{m+1}=1}^{I_{m+1}} \dots \sum_{i_M=1}^{I_M} K(\mathbf{x}_{i_1 \dots i_{m-1} j i_{m+1} \dots i_M}, \mathbf{x}_{i_1 \dots i_{m-1} k i_{m+1} \dots i_M}). \quad (11)$$

- (c) $\left\{ \begin{array}{l} \text{For K-MPCA: Set } \mathbf{C}_m := \mathbf{U}, \text{ the left matrix of the SVD of } [\mathbf{X}_{[m]} \mathbf{X}_{[m]}^T] \text{ from (11)} \\ \text{Truncate to } \tilde{R}_m \text{ columns } \mathbf{U}_m \in \mathbb{C}^{I_m \times \tilde{R}_m}. \\ \text{For K-MICA: Set } \mathbf{C}_m := \mathbf{U}_m \mathbf{W}_m^{-1}. \text{ The additional rotation matrix } \mathbf{W}_m \text{ may be} \\ \text{computed based on negentropy, mutual information, or higher-} \\ \text{order cumulants [85]. The initial SVD of } [\mathbf{X}_{[m]} \mathbf{X}_{[m]}^T] \text{ from (11)} \\ \text{truncates the subspace to } \tilde{R}_m. \end{array} \right.$

2. Set $\mathcal{T} := \mathcal{X}_M \times_M \mathbf{C}_M^+$. For K-MPCA, $\mathbf{C}_M^+ = \mathbf{C}_M^T$.

Output the converged extended core tensor $\mathcal{T} \in \mathbb{C}^{I_0 \times \tilde{R}_1 \times \dots \times \tilde{R}_M}$ and causal factor mode matrices $\mathbf{C}_1, \dots, \mathbf{C}_M$.

^a Every SVD step may be autoencoder-decoder. See Algorithm 1, Footnotes *a* and *b*.
See Fig. 3 for a scalable neural network implementation.

Linear kernel:	$K(\mathbf{u}, \mathbf{v}) = \mathbf{u}^T \mathbf{v} = \mathbf{u} \cdot \mathbf{v}$
Polynomial kernel of degree d :	$K(\mathbf{u}, \mathbf{v}) = (\mathbf{u}^T \mathbf{v})^d$
Polynomial kernel up to degree d :	$K(\mathbf{u}, \mathbf{v}) = (\mathbf{u}^T \mathbf{v} + 1)^d$
Sigmoidal kernel:	$K(\mathbf{u}, \mathbf{v}) = \tanh(\alpha \mathbf{u}^T \mathbf{v} + \beta)$
Gaussian (radial basis function (RBF)) kernel:	$K(\mathbf{u}, \mathbf{v}) = \exp\left(-\frac{\ \mathbf{u} - \mathbf{v}\ ^2}{2\sigma^2}\right)$

Table 1: Common kernel functions. Kernel functions are symmetric, positive semi-definite functions corresponding to symmetric, positive semi-definite Gram matrices. The linear kernel does not modify or warp the feature space.

When a causal factor is a combination of multiple independent sources that are causal in nature, we employ a rotation matrix \mathbf{W} to identify them. The rotation matrix is computed by employing either mutual information, negentropy, or higher-order cumulants [16,5,35,4]. A *Kernel Multilinear ICA (K-MICA) Algorithm* is a kernel generalization of the multilinear independent component analysis (MICA) algorithm [85]. Algorithm 2 simultaneously specifies both K-MPCA and K-MICA algorithms. A scalable tensor factor analysis represents an observation as a hierarchy of parts and wholes [80,79].

2.4 Neural Network Architecture

Tensor factor analysis models are transformed into causal neural networks by using autoencoder-decoders as building blocks. Causal neural networks are composed of causal capsules and tensor transformers, Fig. 1. *Causal capsules* estimate a set of latent variables that represent the causal factors of data formation. A *tensor transformer* governs the causal factor interaction. The M-mode SVD (Algorithm 1) is transformed into a neural network by replacing every SVD step with gradient descent optimization, which is outsourced to an autoencoder-decoder with neurons that have a linear transfer function, Supplemental A. For effectiveness, we employ stochastic gradient descent [9,62]. The extended core tensor $\mathbf{T}_{[0]}$ is computed by defining and employing a tensor autoencoder, an autoencoder whose code is initialized to the Kronecker product, \otimes , of the causal factor representations, $\{\mathbf{u}_{i_m} | 1 \leq i_m \leq I_m \text{ and } 1 \leq m \leq M\}$,

$$\mathbf{d}_{i_1, \dots, i_m, \dots, i_M} = \mathcal{T} \times_1 \hat{\mathbf{u}}_1^T \times \dots \times \hat{\mathbf{u}}_m^T \times \dots \times \hat{\mathbf{u}}_M^T \quad (12)$$

$$= \mathbf{T}_{[0]} (\mathbf{u}_{i_M}^T \otimes \dots \otimes \mathbf{u}_{i_m}^T \otimes \dots \otimes \mathbf{u}_{i_1}^T)^T. \quad (13)$$

To address a set of arbitrarily non-linear causal factors, each autoencoder employs kernel functions (Algorithm 2, Table 1).

2.5 Causal Deep Networks and Scalable Tensor Factor Analysis:

For a scalable architecture, we leverage the properties of block algebra. Shallow autoencoders are replaced with either a mathematically equivalent deep neural network that is a part-based hierarchy of autoencoders-decoders (*i.e.*, feed-forward of Restricted Boltzman Machines [45]), or a set of concurrent autoencoders-decoders, Fig. 3.

For example, the orthonormal subspace of a data batch, $\mathbf{D} \in \mathbb{C}^{I_0 \times I_1}$ that has I_0 measurements and I_1 observations may be computed by recursively subdividing the data, analyzing the data blocks, dimensionally reducing their representations, and finally merging the information

$$\begin{aligned} \mathbf{D} &= \begin{bmatrix} \mathbf{D}_A \\ \mathbf{D}_B \end{bmatrix} = \begin{bmatrix} \mathbf{U}_A \mathbf{S}_A \mathbf{V}_A^T \\ \mathbf{U}_B \mathbf{S}_B \mathbf{V}_B^T \end{bmatrix} = \begin{bmatrix} \mathbf{U}_A & \mathbf{0} \\ \mathbf{0} & \mathbf{U}_B \end{bmatrix} \underbrace{\begin{bmatrix} \mathbf{S}_A \mathbf{V}_A^T \\ \mathbf{S}_B \mathbf{V}_B^T \end{bmatrix}}_{\text{SVD}} = \begin{bmatrix} \mathbf{U}_A & \mathbf{0} \\ \mathbf{0} & \mathbf{U}_B \end{bmatrix} \mathbf{W} \Sigma \mathbf{V}^T \\ &= \begin{bmatrix} \mathbf{U}_A & \mathbf{0} \\ \mathbf{0} & \mathbf{U}_B \end{bmatrix} \begin{bmatrix} \mathbf{W}_{AA} & \mathbf{W}_{AB} \\ \mathbf{W}_{BA} & \mathbf{W}_{BB} \end{bmatrix} \Sigma \mathbf{V}^T = \begin{bmatrix} \mathbf{U}_A \mathbf{W}_{AA} \\ \mathbf{U}_B \mathbf{W}_{BB} \end{bmatrix} \Sigma \mathbf{V}^T = \mathbf{U} \Sigma \mathbf{V}^T. \end{aligned} \quad (14)$$

The rotation matrix \mathbf{W} transforms basis matrices, \mathbf{U}_A and \mathbf{U}_B , which span the subspaces of the observations in the data blocks \mathbf{D}_A and \mathbf{D}_B , such that the observations in \mathbf{D}_A and the corresponding observations in \mathbf{D}_B no longer have distinct representations \mathbf{V}_A^T and \mathbf{V}_B^T , but have the same representations \mathbf{V}^T . The representations \mathbf{V}_A^T and \mathbf{V}_B^T may be dimensionally reduced before moving onto the next step.⁶

Computing causal factor representations, the mode matrices \mathbf{U}_m of an MPCA tensor model, is equivalent to computing a M different of mutually constrained, cluster-based

⁶ Block algebra may be employed if the tensor model is multilinear (tensor) principal component analysis (MPCA), multilinear (tensor) independent component analysis (MICA)[85], Kernel-MPCA or Kernel-MICA[77].

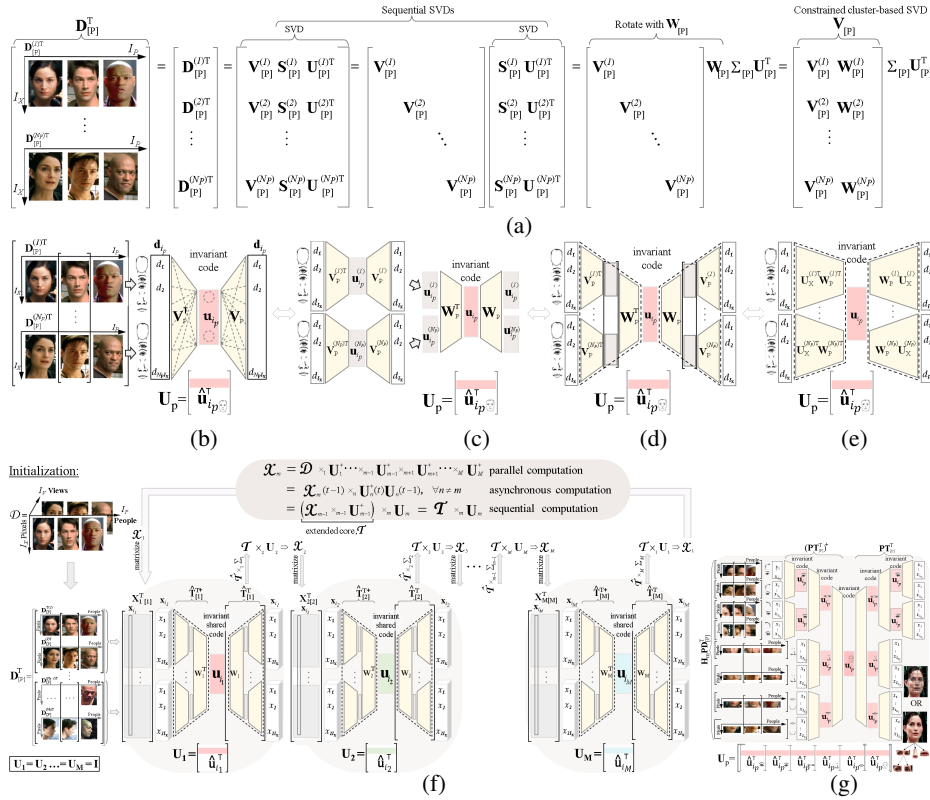


Fig. 3: Deep neural network. Subfigures (a-e) depict [14–19] and [77, pg.38-40].

(a) The mode matrix computation \mathbf{U}_m may be thought as a constrained cluster-based PCA that is rewritten in terms of block SVDs. Matrixizing \mathcal{D} may be viewed as a concatenation of “cluster” data. The matrix \mathbf{W} transforms the basis matrix $\mathbf{V}_0^{(n)}$ such that the causal factor representation \mathbf{U}_m is the same regardless of cluster membership. In a tensor model, there are M different constrained cluster-based PCAs. (b) Mode matrix \mathbf{U}_m computation using a single autoencoder-decoder. (c) Mode matrix computation as a hierarchy of autoencoder-decoders (*i.e.*, feed-forward [45]), (d) Mode matrix computation written as a deep learning model (e) Concurrent-autoencoders; *i.e.*, constrained cluster-based autoencoders. (f) Forward causal model with a set of capsules implemented by deep neural networks. For a parallel, synchronized or asynchronous computation, we break the chain links and shuttle causal information, \mathbf{U}_m , between capsules to compute $\mathcal{X}(t+1)$ for the next iteration. (g) Each capsule in (f) may be replaced with a part-based deep neural network by permuting the rows in $\mathbf{D}_{[m]}^T$ with \mathbf{P} , segmented by \mathbf{H}_m , which is efficiently trained with a part-based hierarchy of autoencoders, Fig. 4.

PCA, [77, pg.38-40], Fig. 3a. When dealing with data that can be separated into clusters, the standard machine learning approach is to compute a separate PCA. When data from different clusters are generated by the same underlying process (e.g., facial images of the same people under different viewing conditions), the data blocks can be concatenated in the measurement mode and the common causal factor can be modeled by one PCA.




Fig. 4: (a) Causal capsules may be implemented with a part-based hierarchy of autoencoders. The dataset is permuted by \mathbf{P} , segmented and filtered by \mathbf{H}_m that is mode dependent. (b) Implementing the capsules with a part-based hierarchy of autoencoders is equivalent to performing M-mode Block SVD[80, Sec IV][76].

However, for a scalable solution, we employ block algebra (14) and compute a set of constrained cluster-based PCAs, *i.e.*, a set of concurrent PCAs. Thus, we define a *constrained, cluster-based PCA* as the computation of a set of PCA basis vectors, such that the latent representation is constrained to be the invariant of the cluster membership.

In the context of our multifactor data analysis, we define a cluster as a set of observations for which all factors are fixed but one. For every tensor mode, there are $N_m = I_1 I_2 \dots I_{m-1} I_{m+1} \dots I_M$ possible clusters and the data in each cluster varies with the same causal mode. The constrained, cluster-based PCA concatenates the clusters in the measurement mode and analyzes the data with a linear model, such as PCA.

To see this, let $\mathcal{D}_{i_1 \dots i_{m-1} i_{m+1} \dots i_M} \in \mathbb{C}^{I_0 \times 1 \times 1 \dots \times 1 \times I_m \times 1 \times 1 \dots \times 1}$ denote a subtensor of \mathcal{D} that is obtained by fixing all causal factor modes but mode m and mode 0 (the measurement mode). Matrixizing this subtensor in the measurement mode we obtain $\mathbf{D}_{i_1 \dots i_{m-1} i_{m+1} \dots i_M [0]} \in \mathbb{C}^{I_0 \times I_m}$. This data matrix comprises a cluster of data obtained by varying causal factor m , to which one can traditionally apply PCA. Since there are $N_m = I_1 I_2 \dots I_{m-1} I_{m+1} \dots I_M$ possible clusters that share the same underlying space associated with factor m , the data can be concatenated and PCA performed in order to extract the same representation for factor m regardless of the cluster. Now, consider the MPCA computation of mode matrix \mathbf{U}_m , Fig. 3a, which can be written in terms of matrixized subtensors as

$$\mathbf{D}_m = \begin{bmatrix} \mathbf{D}_{1 \dots 1 1 \dots 1 [m]}^T \\ \vdots \\ \mathbf{D}_{I_1 \dots 1 1 \dots 1 [m]}^T \\ \vdots \\ \mathbf{D}_{I_1 \dots I_{m-1} I_{m+1} \dots I_M [m]}^T \end{bmatrix}^T = \mathbf{U}_m \mathbf{\Sigma}_m \mathbf{V}_m^T. \quad (15)$$

This is equivalent to computing a set of $N_m = I_1 I_2 \dots I_{m-1} I_{m+1} \dots I_M$ cluster-based PCAs concurrently by combining them into a single statistical model and representing the underlying causal factor m common to the clusters. Thus, rather than computing a separate linear PCA model for each cluster, MPCA concatenates the clusters into a single statistical model and computes a representation (coefficient vector) for mode m that is invariant relative to the other causal factor modes $1, \dots, (m-1), (m+1), \dots, M$.

For a scalable solution, we rotate the cluster-based PCA basis vectors, such that the data blocks have the same representation regardless of cluster membership. Thus, MPCA performs multiple constrained, cluster-based PCA. To clarify the relationship, let us number each of the matrices $\mathbf{D}_{i_1 \dots i_{m-1} i_{m+1} \dots i_M} = \mathbf{D}_m^{(n)}$ with a parenthetical superscript $1 \leq n = 1 + \sum_{k=1, k \neq m}^M (i_k - 1) \prod_{l=1, l \neq m}^{k-1} I_l \leq N_m$. Let each of the cluster SVDs be $\mathbf{D}_m^{(n)} = \mathbf{U}_m^{(n)} \Sigma_m^{(n)} \mathbf{V}_m^{(n)\top}$, and

$$\mathbf{D}_{[m]} = \underbrace{[\mathbf{U}_m^{(1)} \Sigma_m^{(1)} \dots \mathbf{U}_m^{(N_m)} \Sigma_m^{(N_m)}]}_{\text{SVD}} \text{diag}([\mathbf{V}_m^{(1)} \dots \mathbf{V}_m^{(N_m)}])^\top \quad (16)$$

$$= \mathbf{U}_m \Sigma_m \mathbf{W}_m^\top \text{diag}([\mathbf{V}_m^{(1)} \dots \mathbf{V}_m^{(N_m)}])^\top, \quad (17)$$

$$= \mathbf{U}_m \Sigma_m [\mathbf{V}_m^{(1)} \mathbf{W}_m^{(1)} \dots \mathbf{V}_m^{(N_m)} \mathbf{W}_m^{(N_m)}]^\top \quad (18)$$

$$= \mathbf{U}_m \Sigma_m \mathbf{V}_m^\top, \quad (19)$$

where $\text{diag}(\cdot)$ denotes a diagonal matrix whose elements are each of the elements of its vector argument. The mode matrix $\mathbf{V}_m^{(n_m)}$ is the measurement matrix $\mathbf{U}_0^{(n_m)}$ ($\mathbf{U}_x^{(n_m)}$ when the measurements are image pixels) that contains the eigenvectors spanning the observed data in cluster n_m , $1 \leq n_m \leq N_m$. MPCA can be thought as computing a rotation matrix, \mathbf{W}_m , that contains a set of blocks $\mathbf{W}_m^{(n)}$ along the diagonal that transform the PCA cluster eigenvectors $\mathbf{V}_m^{(n_m)}$ such that the mode matrix \mathbf{U}_m is the same regardless of cluster membership (16–19), Fig 3. The constrained “cluster”-based PCAs may also be implemented with a set of concurrent “cluster”-based PCAs, Fig. 3e.

Causal factors of object wholes may be computed efficiently from their parts, by applying a permutation matrix \mathbf{P} and creating part-based data clusters with a segmentation filter \mathbf{H}_m , where $\mathbf{D}^\top \times_m \mathbf{H}_m \mathbf{P} \Leftrightarrow \mathbf{H}_m \mathbf{P} \mathbf{D}_{[m]}^\top$. A deep neural network can be efficiently trained with a hierarchy of part-based autoencoders, Fig. 4. A computation that employs a part-based hierarchy of autoencoders parallels the Incremental M-mode Block SVD [80, 76, Sec. IV]. A data tensor is recursively subdivided into data blocks, analyzed in a bottom-up fashion, and the results merged as one moves through the hierarchy. The computational cost is the cost of training one autoencoder, $\mathcal{O}(T)$, times $\mathcal{O}(\log N_m)$, the total number of autoencoders trained for each factor matrix, $\mathcal{O}(T \log N_m)$. If the causal neural network is trained sequentially, the training cost for one time iteration is $\mathcal{O}(MT \log \bar{N})$, where \bar{N} is the average number of clusters across the M modes.

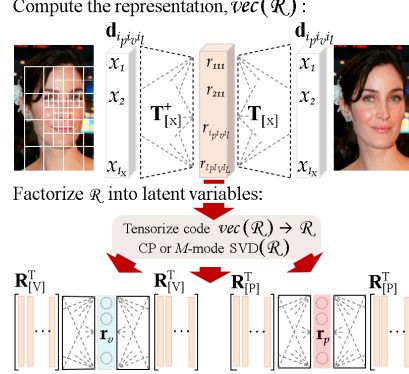
3 Inverse Causal Question: “Why?”

Inverse causal inference addresses the “why” question and estimates the causes of effects given an estimated forward causal model and a set of constraints that reduce the solution set⁷ and render the problem well-posed [26, 78, 86].

Multilinear tensor factor analysis constrains causal factor representations to be unitary vectors. Multilinear projection [86, 78] relies on this constraint and performs multiple regularized regressions. One or more unlabeled test observations that are not part of the training data set are multilinearly projected into the causal factor spaces,

⁷ As discussed earlier in the paper, different combinations of the same causal factors can lead to the same outcome. In imaging, this phenomena may result in visual illusions. This is a many-to-one problem, and its inverse is ill posed without constraints.

Fig. 5: An inverse causal network is an inverted (upside-down) forward network that performs its operations in reverse order using $\mathbf{T}_{[X]}^+$, estimated during training, Fig. 1. The architecture implements the multilinear projection [78, 86]. For a scalable solution, each autoencoder-decoder is replaced with a deep network.



$$\begin{aligned} \mathcal{T}^{\dagger_0} \times_0^T \mathbf{d}_{\text{test}} &= \mathcal{R} \\ &= \text{CP}(\mathcal{R}) \text{ or } \text{M-mode SVD}(\mathcal{R}) \\ &\approx \hat{\mathbf{r}}_1 \dots \circ \hat{\mathbf{r}}_m \dots \circ \hat{\mathbf{r}}_M, \end{aligned} \quad (20)$$

where $\hat{\mathbf{r}}_m$ is the estimated orthonormal causal factor m representation, \circ is the outer product and \mathcal{T}^{\dagger_0} is the tensorized $\mathbf{T}_{[0]}^+$. A neural network that implements a multilinear projection architecture is an inverted (upside down) forward neural network architecture that employs an estimated $\mathbf{T}_{[X]}^+$ and reverses the operation order, Fig. 5.

Neural architectures addressing underdetermined inverse problems are characterized by hidden layers that are wider than the input layer; *i.e.*, the dimensionality of $vec(\mathcal{R})$ is larger than the number of measurements in \mathbf{d} . Dimensionality reduction reduces noise, and the width of the hidden layers [30]. However, it can also camouflage an inherently underdetermined inverse problem. Adding sparsity, non-negativity constraints [87], etc., can further reduce the solution set in a principled way. Alternatively or in addition, one can determine a set of candidate solutions by modeling different aspects of the mechanism of data formation as piecewise tensor (multilinear) factor models. A single multilinear projection [86, 78] is replaced with multiple multilinear projections. Vasilescu and Terzopoulos [81] rewrote the forward multilinear model in terms of multiple piecewise linear models that were employed to perform multiple linear projections and produced multiple candidate solutions that were gated to return the most likely solution.

4 Conclusion

We derive a set of shallow and deep causal neural networks that are a consequence of causal tensor factor analysis. Causal neural networks are composed of causal capsules and a tensor transformer. Causal capsules compute invariant causal factor representations, whose interaction are governed by a tensor transformation. An inverse causal neural network implements the multilinear projection and estimates the causes of effects. As an alternative to aggressive “bottleneck” dimensionality reduction that may camouflage an inherently underdetermined inverse problem, the mechanism of data formation is modeled as piecewise tensor (multilinear) models, and inverse causal neural networks perform multiple multilinear projections that result in multiple candidate solutions, which may be gated to yield the most likely solution.

Acknowledgement: Demetri Terzopoulos and Ernest Davis provided invaluable feedback on various drafts of this paper.

Supplemental

- A. **PCA computation with a Hebb Autoencoder**
- B. **Relevant Tensor Algebra**
- C. **Compositional Hierarchical Block TensorFaces**

References

1. D. H. Ackley, G. A. Hinton, and T. J. Sejnowski. A learning algorithm for Boltzmann machines. *Cognitive Science*, 9(1):147–169, 1985.
2. J. D. Angrist. Lifetime earnings and the vietnam era draft lottery: evidence from social security administrative records. *The American Economic Review*, pages 313–336, 1990.
3. F. R. Bach and M. I. Jordan. Kernel independent component analysis. *Journal of Machine Learning Research*, 3(Jul):1–48, 2002.
4. M. Bartlett, J. Movellan, and T. Sejnowski. Face recognition by independent component analysis. *IEEE Transactions on Neural Networks*, 13(6):1450–64, 2002.
5. A. J. Bell and T. J. Sejnowski. An information-maximization approach to blind separation and blind deconvolution. *Neural Computation*, (6):1004–1034, 1995.
6. J. Benesty, C. Paleologu, L. Dogariu, and S. Ciochină. Identification of linear and bilinear systems: A unified study. *Electronics*, 10(15), 2021.
7. Y. Bengio and A. Courville. *Handbook on Neural Information Processing*, chapter Deep Learning of Representations, pages 1–28. Springer Berlin Heidelberg, 2013.
8. P. Bentler and S. Lee. A statistical development of three-mode factor analysis. *British J. of Math. and Stat. Psych.*, 32(1):87–104, 1979.
9. L. Bottou et al. Online learning and stochastic approximations. *On-line learning in neural networks*, 17(9):142, 1998.
10. R. Bro. Parafac: Tutorial and applications. In *Chemom. Intell. Lab Syst., Special Issue 2nd Internet Cont. in Chemometrics (INCINC’96)*, volume 38, pages 149–171, 1997.
11. A. Bulat, J. Kossaifi, G. Tzimiropoulos, and M. Pantic. Incremental multi-domain learning with network latent tensor factorization. In *The Thirty-Fourth AAAI Conference on Artificial Intelligence, AAAI*, pages 10470–10477. AAAI Press, 2020.
12. D. Card and A. B. Krueger. Minimum wages and employment: A case study of the fast food industry in New Jersey and Pennsylvania. *National Bureau of Economic Research Cambridge, Mass., USA*, 1993.
13. J. D. Carroll and J. J. Chang. Analysis of individual differences in multidimensional scaling via an N-way generalization of ‘Eckart-Young’ decomposition. *Psychometrika*, 35:283–319, 1970.
14. J. C. Chen, R. Ranjan, A. Kumar, C. H. Chen, V. M. Patel, and R. Chellappa. An end-to-end system for unconstrained face verification with deep convolutional neural networks. In *IEEE International Conf. on Computer Vision Workshop (ICCVW)*, pages 360–368, Dec 2015.
15. W. Chu and Z. Ghahramani. Probabilistic models for incomplete multi-dimensional arrays. volume 5 of *Proc. of Machine Learning Research(PMLR)*, pages 89–96, Apr 2009.
16. P. Common. Independent component analysis, a new concept? *Signal Processing*, 36:287–314, 1994.
17. J. Davis and H. Gao. Recognizing human action efforts: An adaptive three-mode PCA framework. In *Proc. IEEE Inter. Conf. on Computer Vision, (ICCV)*, pages 1463–69, Nice, France, Oct 13-16 2003.
18. L. de Lathauwer, B. de Moor, and J. Vandewalle. A multilinear singular value decomposition. *SIAM J. of Matrix Analysis and Applications*, 21(4):1253–78, 2000.
19. L. de Lathauwer, B. de Moor, and J. Vandewalle. On the best rank-1 and rank- (R_1, R_2, \dots, R_n) approximation of higher-order tensors. *SIAM J. of Matrix Analysis and Applications*, 21(4):1324–42, 2000.
20. A. P. Dempster, N. M. Laird, and D. B. Rubin. Maximum likelihood from incomplete data via the em algorithm. *Journal of the Royal Statistical Society: Series B (Methodological)*, 39(1):1–22, 1977.
21. G. Desjardins, A. Courville, and Y. Bengio. Disentangling factors of variation via generative entangling. *arXiv:1210.5474*, 2012.
22. J. M. Dudek, L. Dueñas-Osorio, and M. Y. Vardi. Efficient contraction of large tensor networks for weighted model counting through graph decompositions, 2019.
23. H. Fan, B. Xiong, K. Mangalam, Y. Li, Z. Yan, J. Malik, and C. Feichtenhofer. Multiscale vision transformers. In *Proc. of the IEEE/CVF Inter. Conf. on Computer Vision*, pages 6824–35, 2021.

24. B. C. v. Fraassen. *The Scientific Image*. Oxford University Press, 12 1980.
25. T. Gebru, J. Morgenstern, B. Vecchione, J. W. Vaughan, H. Wallach, H. Daumé III, and K. Crawford. Datasheets for datasets. *Commun. ACM*, 64(12):86–92, nov 2021.
26. A. Gelman and G. Imbens. Why ask why? Forward causal inference and reverse causal questions. Tech.report, Nat.Bureau of Econ. Research, 2013.
27. D. B. Grimes and R. P. Rao. Bilinear sparse coding for invariant vision. *Neural computation*, 17(1):47–73, 2005.
28. G. Grindlay and M. A. O. Vasilescu. A multilinear (tensor) framework for hrtf analysis and synthesis. In 2007 IEEE Inter. Conf. on Acoustics, Speech and Signal Processing - ICASSP '07, volume 1, pages I–161–164, 2007.
29. R. Harshman. Foundations of the PARAFAC procedure: Model and conditions for an explanatory factor analysis. Tech. Report Working Papers in Phonetics 16, UCLA, CA, Dec 1970.
30. G. E. Hinton and R. R. Salakhutdinov. Reducing the dimensionality of data with neural networks. *Science*, 313(5786):504–507, 2006.
31. P. W. Holland. Statistics and causal inference: Rejoinder. *J. of the American Statistical Association*, 81(396):968–970, 1986.
32. R. C. Hoover, K. Caudle, and K. Braman. A new approach to multilinear dynamical systems and control, 2021.
33. G. B. Huang. Learning hierarchical representations for face verification with convolutional deep belief networks. In IEEE Conf. on Computer Vision and Pattern Recognition (CVPR), pages 2518–25, Jun 2012.
34. H. Wang and N. Ahuja. Facial expression decomposition. In Proc, 9th IEEE Inter. Conf. on Computer Vision (ICCV), pages 958–65, v.2, 2003.
35. A. Hyvärinen, J. Karhunen, and E. Oja. *Independent Component Analysis*. Wiley, NY, 2001.
36. G. Imbens and D. Rubin. *Causal Inference for Statistics, Social and Biomedical Sciences: An Introduction*. Cambridge Univ. Press, 2015.
37. G. W. Imbens and J. D. Angrist. Identification and estimation of local average treatment effects. *Econometrica*, 62(2):467–475, 1994.
38. M. A. Iwen, D. Needell, E. Rebrova, and A. Zare. Lower memory oblivious (tensor) subspace embeddings with fewer random bits: modewise methods for least squares. *SIAM Journal on Matrix Analysis and Applications*, 42(1):376–416, 2021.
39. A. Kapteyn, H. Neudecker, and T. Wansbeek. An approach to n -mode component analysis. *Psychometrika*, 51(2):269–275, Jun 1986.
40. D. S. Kermany, M. Goldbaum, W. Cai, C. C. Valentim, H. Liang, S. L. Baxter, A. McKeown, G. Yang, X. Wu, F. Yan, J. Dong, M. K. Prasadha, J. Pei, M. Y. Ting, J. Zhu, C. Li, S. Hewett, J. Dong, I. Ziyar, A. Shi, R. Zhang, L. Zheng, R. Hou, W. Shi, X. Fu, Y. Duan, V. A. Huu, C. Wen, E. D. Zhang, C. L. Zhang, O. Li, X. Wang, M. A. Singer, X. Sun, J. Xu, A. Tafreshi, M. A. Lewis, H. Xia, and K. Zhang. Identifying medical diagnoses and treatable diseases by image-based deep learning. *Cell*, 172(5):1122–1131.e9, 2018.
41. V. Khrulkov. *Geometrical Methods in Machine Learning and Tensor Analysis*. PhD dissertation, Skolkovo Institute, 2020.
42. Y. Kim, E. Park, S. Yoo, T. Choi, L. Yang, and D. Shin. Compression of deep convolutional neural networks for fast and low power mobile applications. *CoRR*, abs/1511.06530, 2015.
43. A. Krizhevsky, I. Sutskever, and G. E. Hinton. Imagenet classification with deep convolutional neural networks. In *Neural Information Processing Systems 25 (NIPS)*, volume 25, pages 1097–1105. Curran Associates, Inc., 2012.
44. P. M. Kroonenberg and J. de Leeuw. Principal component analysis of three-mode data by means of alternating least squares algorithms. *Psychometrika*, 45:69–97, 1980.
45. H. Larochelle, D. Erhan, A. Courville, J. Bergstra, and Y. Bengio. An empirical evaluation of deep architectures on problems with many factors of variation. In *Proceedings of the 24th International Conference on Machine Learning*, page 473–480, New York, NY, USA, 2007.
46. V. Lebedev, Y. Ganin, M. Rakhuba, I. V. Oseledets, and V. S. Lempitsky. Speeding-up convolutional neural networks using fine-tuned cp-decomposition. *CoRR*, abs/1412.6553, 2014.

47. D. Lewis. Causal Explanation. In Philosophical Papers Volume II, pages 214–240. Oxford University Press, 06 1987.
48. Z. C. Lipton. The Mythos of Model Interpretability. Commun. ACM, 61(10):36–43, 2018.
49. Z. C. Lipton and J. Steinhardt. Troubling trends in machine learning scholarship: Some ml papers suffer from flaws that could mislead the public and stymie future research. Queue, 17(1):45–77, Feb 2019.
50. Z. Liu, Y. Lin, Y. Cao, H. Hu, Y. Wei, Z. Zhang, S. Lin, and B. Guo. Swin transformer: Hierarchical vision transformer using shifted windows. In Proceedings of the IEEE/CVF International Conference on Computer Vision, pages 10012–10022, 2021.
51. A. Madani, M. Moradi, A. Karargyris, and T. Syeda-Mahmood. Semi-supervised learning with generative adversarial networks for chest x-ray classification with ability of data domain adaptation. In 2018 IEEE 15th Inter. Symp. on Biomedical Imaging (ISBI 2018), pages 1038–1042, 2018.
52. J. Magnus and H. Neudecker. Matrix Differential Calculus with Applications in Statistics and Econometrics. John Wiley & Sons, 1988.
53. M. F. Mathieu, J. J. Zhao, J. Zhao, A. Ramesh, P. Sprechmann, and Y. LeCun. Disentangling factors of variation in deep representation using adversarial training. Adv. in Neural Information Processing Systems, 29, 2016.
54. R. Memisevic and G. E. Hinton. Learning to Represent Spatial Transformations with Factored Higher-Order Boltzmann Machines. Neural Computation, 22(6):1473–1492, 06 2010.
55. A. Novikov, D. Podoprikin, A. Osokin, and D. P. Vetrov. Tensorizing neural networks. In Neural Information Processing Systems 28, pages 442–450. Curran Associates, Inc., 2015.
56. E. Oja. A simplified neuron model as a principal component analyzer. Journal of Mathematical Biology, 15:267–2735, 1982.
57. C. C. Onu, J. E. Miller, and D. Precup. A fully tensorized recurrent neural network. CoRR, abs/2010.04196, 2020.
58. J. Pearl. Causality: Models, Reasoning, and Inference. Cambridge Univ. Press, 2000.
59. J. Pearl and E. Bareinboim. External validity: From do-calculus to transportability across populations. Statistical Science, 29(4):579–95, 2014.
60. R. Penrose. Applications of negative-dimensional tensors. In D. Welsh, editor, Combinatorial Mathematics and Its Applications. Academic Press, 1971.
61. Q. Qiu and R. Chellappa. Compositional dictionaries for domain adaptive face recognition. IEEE Transactions on Image Processing, 24(12):5152–5165, 2015.
62. H. Robbins and S. Monro. A stochastic approximation method. The annals of mathematical statistics, pages 400–407, 1951.
63. D. B. Rubin. For objective causal inference, design trumps experimental analysis. The Annals of Applied Statistics, 2(3):808 – 840, 2008.
64. D. E. Rumelhart, G. E. Hinton, and R. J. Williams. Learning internal representations by error propagation. Parallel Distributed Processing, 1986.
65. T. Sanger. Optimal unsupervised learning in a single layer linear feedforward neural network. Neural Networks, 12:459–473, 1989.
66. B. Schölkoph, A. Smola, and K.-R. Muller. Nonlinear component analysis as a kernel eigenvalue problem. Neural Computation, 10(5):1299–1319, 1998.
67. T. Sejnowski, S. Chattarji, and P. Sfanton. Induction of Synaptic Plasticity by Hebbian Covariance in the Hippocampus, pages 105–124. Addison-Wesley, 1989.
68. P. Spirtes, C. N. Glymour, R. Scheines, and D. Heckerman. Causation, prediction, and search. MIT press, 2000.
69. Y. Sun, X. Wang, and X. Tang. Hybrid deep learning for face verification. In Proc. IEEE International Conf. on Computer Vision (ICCV), pages 1489–96, Dec 2013.
70. Y. Taigman, M. Yang, M. Ranzato, and L. Wolf. Deepface: Closing the gap to human-level performance in face verification. In Proc. IEEE Conf. on Computer Vision and Pattern Recognition, pages 1701–08, 2014.
71. Y. Tang, R. Salakhutdinov, and G. Hinton. Tensor analyzers. volume 28 of Proceedings of Machine Learning Research, pages 163–171, Atlanta, Georgia, USA, 17–19 Jun 2013.

72. N. Tishby and N. Zaslavsky. Deep learning and the information bottleneck principle. In 2015 IEEE Information Theory Workshop (ITW), pages 1–5, 2015.
73. E. J. Topol. High-performance medicine: the convergence of human and artificial intelligence. Nature Medicine, 25(1):44–56, 2019.
74. L. R. Tucker. Some mathematical notes on three-mode factor analysis. Psychometrika, 31:279–311, 1966.
75. M. A. O. Vasilescu. Human motion signatures: Analysis, Synthesis, Recognition. In Proc. Int. Conf. on Pattern Recognition, volume 3, pages 456–460, Quebec City, Aug 2002.
76. M. A. O. Vasilescu. Incremental Multilinear SVD. In Proc. Conf. on ThRee-way methods In Chemistry And Psychology (TRICAP 06), 2006.
77. M. A. O. Vasilescu. A Multilinear (Tensor) Algebraic Framework for Computer Graphics, Computer Vision, and Machine Learning. PhD dissertation, University of Toronto, 2009.
78. M. A. O. Vasilescu. Multilinear projection for face recognition via canonical decomposition. In Proc. IEEE Inter. Conf. on Automatic Face Gesture Recognition (FG 2011), pages 476–483, Mar 2011.
79. M. A. O. Vasilescu and E. Kim. Compositional hierarchical tensor factorization: Representing hierarchical intrinsic and extrinsic causal factors. In The 25th ACM SIGKDD Conf. on Knowledge Discovery and Data Mining (KDD 2019): Tensor Methods for Emerging Data Science Challenges Workshop, Aug. 5 2019.
80. M. A. O. Vasilescu, E. Kim, and X. S. Zeng. CausalX: Causal eXplanations and block multilinear factor analysis. In 2020 25th International Conference of Pattern Recognition (ICPR 2020), pages 10736–10743, Jan 2021.
81. M. A. O. Vasilescu and D. Terzopoulos. Multilinear analysis for facial image recognition. In Proc. Int. Conf. on Pattern Recognition, volume 2, pages 511–514, Quebec City, Aug 2002.
82. M. A. O. Vasilescu and D. Terzopoulos. Multilinear analysis of image ensembles: TensorFaces. In Proc. European Conf. on Computer Vision (ECCV 2002), pages 447–460, May 2002.
83. M. A. O. Vasilescu and D. Terzopoulos. Multilinear subspace analysis of image ensembles. In Proc. IEEE Conf. on Computer Vision and Pattern Recognition, volume II, pages 93–99, 2003.
84. M. A. O. Vasilescu and D. Terzopoulos. TensorTextures: Multilinear Image-Based Rendering. ACM Transactions on Graphics, 23(3):336–342, Aug 2004. Proc. ACM SIGGRAPH Conf.
85. M. A. O. Vasilescu and D. Terzopoulos. Multilinear independent components analysis. In Proc. IEEE Conf. on Computer Vision and Pattern Recognition, pages 547–553, v.I, 2005.
86. M. A. O. Vasilescu and D. Terzopoulos. Multilinear projection for appearance-based recognition in the tensor framework. In Proc. 11th IEEE Inter. Conf. on Computer Vision (ICCV’07), pages 1–8, 2007.
87. J. Vendrow, J. Haddock, and D. Needell. A generalized hierarchical nonnegative tensor decomposition, 2021.
88. J.-P. Vert, K. Tsuda, and B. Schölkopf. A primer on kernel methods. Kernel methods in computational biology, 47:35–70, 2004.
89. D. Vlasic, M. Brand, H. Pfister, and J. Popovic. Face transfer with multilinear models. ACM Transactions on Graphics (TOG), 24(3):426–433, Jul 2005.
90. H. Wang and N. Ahuja. A tensor approximation approach to dimensionality reduction. Inter. J. of Computer Vision, 6(3):217–29, Mar 2008.
91. W. Wang, E. Xie, X. Li, D.-P. Fan, K. Song, D. Liang, T. Lu, P. Luo, and L. Shao. Pyramid vision transformer: A versatile backbone for dense prediction without convolutions. In Proc. of the IEEE/CVF International Conference on Computer Vision, pages 568–578, 2021.
92. J. Woodward. Making things happen: A theory of causal explanation. Oxford University Press, 2005.
93. C. Xiong, L. Liu, X. Zhao, S. Yan, and T. K. Kim. Convolutional fusion network for face verification in the wild. IEEE Trans. on Circuits and Systems for Video Technology, 26(3):517–28, Mar 2016.
94. J. Yang, X. Gao, D. Zhang, and J. Yang. Kernel ICA: An alternative formulation and its application to face recognition. Pattern Recognition, 38(10):1784–87, 2005.
95. J. Ye. Generalized low rank approximations of matrices. Machine Learning, 61(1):167–191, 2005.