



Identifiability of latent-variable and structural-equation models: from linear to nonlinear

Aapo Hyvärinen¹ · Ilyes Khemakhem² · Ricardo Monti²

Received: 27 January 2023 / Revised: 4 April 2023 / Accepted: 24 April 2023 /
Published online: 4 November 2023

© The Institute of Statistical Mathematics, Tokyo 2023

Abstract

An old problem in multivariate statistics is that linear Gaussian models are often unidentifiable. In factor analysis, an orthogonal rotation of the factors is unidentifiable, while in linear regression, the direction of effect cannot be identified. For such linear models, non-Gaussianity of the (latent) variables has been shown to provide identifiability. In the case of factor analysis, this leads to independent component analysis, while in the case of the direction of effect, non-Gaussian versions of structural equation modeling solve the problem. More recently, we have shown how even general nonparametric nonlinear versions of such models can be estimated. Non-Gaussianity is not enough in this case, but assuming we have time series, or that the distributions are suitably modulated by observed auxiliary variables, the models are identifiable. This paper reviews the identifiability theory for the linear and nonlinear cases, considering both factor analytic and structural equation models.

Keywords Identifiability · Independent component analysis · Structural equation model · Factor analysis · Disentanglement · Non-Gaussianity

1 Introduction

The goal of this paper is to provide a succinct and relatively self-contained exposition of the identifiability theory of a class of latent variable models called independent component analysis, as well as of a class of structural equation models. The

The Related Articles are <https://doi.org/10.1007/s10463-023-00885-3>; <https://doi.org/10.1007/s10463-023-00886-2>; <https://doi.org/10.1007/s10463-023-00887-1>.

✉ Aapo Hyvärinen
aapo.hyvarinen@helsinki.fi

¹ Department of Computer Science, University of Helsinki, Pietari Kalmin katu 5, 00560 Helsinki, Finland

² Gatsby Computational Neuroscience Unit, UCL, 25 Howland Street, London, W1T 4JG, UK

theory has both linear and nonlinear versions, where “nonlinear” is to be taken in the sense of general (nonparametric) nonlinearities. The latent variable models and structural equation model are intimately related, and the identifiability theory of the former can be used to construct an identifiability theory of the latter. We focus on identifiability theory, and aim to explain the basic results on an approachable manner. Estimation methods and algorithms are given very little attention in this paper.

We start by motivating these different models in the rest of this section. The following sections are structured as follows. The notion of identifiability is defined in Sect. 2. The model of (linear) independent component analysis (ICA) is considered in Sect. 3. Linear structural equation models (SEM) are considered in Sect. 4. Moving to nonlinear (nonparametric) models, Sect. 5 treats the identifiability of nonlinear ICA, and Sect. 6, the identifiability of nonlinear SEM. Section 7 provides further discussion on the utility of identifiability, topics for future research, and algorithms. Section 8 concludes the paper.

1.1 Linear representation learning and factor analysis

The problem of identifiability of latent variables was encountered already decades ago in the case of classical factor analysis, which forms the basis for all our developments. The basic model is as follows. Assume $s_i, i = 1, \dots, n$ are n standardized uncorrelated Gaussian latent random variables. The covariance of \mathbf{s} is thus identity, a property which is termed “whiteness.” Assume \mathbf{A} is an $m \times n$ matrix, and denote by \mathbf{n} an m -dimensional vector of uncorrelated noise variables. We observe the m -dimensional random vector \mathbf{x} which is a noisy linear mixture given by

$$\mathbf{x} = \mathbf{A}\mathbf{s} + \mathbf{n}. \quad (1)$$

The goal would be to recover the components in the vector \mathbf{s} , or at least the matrix \mathbf{A} . Many methods have been proposed (Harman 1967), but fundamentally such factor analysis suffers from the indeterminacy of a “factor rotation,” which is another way of saying that the factor analytic model is not identifiable for Gaussian factors. This is because any orthogonal transformation of the Gaussian latent vector, due to its whiteness, gives exactly the same distribution for the data, while giving quite different values of the latent variables, i.e., the features. This is a fundamental problem to which we will return more than once below.

Thus, the parameter matrix \mathbf{A} and the components s_i cannot be uniquely recovered. Such unidentifiability is a problem since one important goal of fitting such models is to find the underlying structure of the data, or a useful *representation*. If the representation a model gives is not unique or even well defined, it is not possible to find the underlying structure of the data. It is therefore crucial to find models which are identifiable. The problem is particularly pertinent in the case where $m = n$, i.e., dimension reduction is not performed, which is our focus in this paper.¹

¹ Dimension reduction in itself does not necessarily suffer from such lack of uniqueness if all we want to find is the right subspace. In the case of dimension reduction, factor analysis is actually closely related to principal component analysis (PCA). However, this whole paper is about the case where dimension reduction is *not* performed; these two are very different problems.

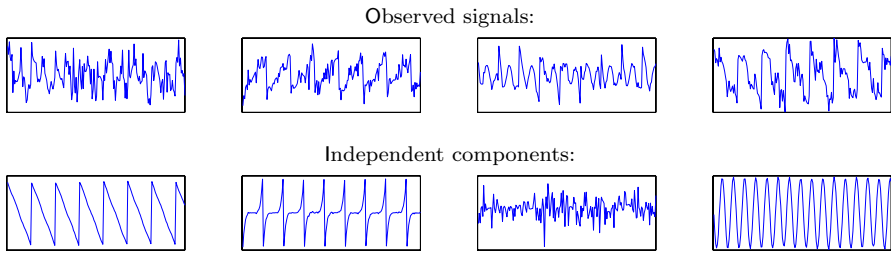


Fig. 1 The basic idea of ICA. From the four measured signals shown in the upper row, ICA is able to recover the original source signals which were mixed together in the measurements, as shown in the bottom row

One practical motivation here is that we might want to separate signals from linear mixtures, and this should be done “blindly,” i.e., using minimum information. This is illustrated in Fig. 1. The four observed signals (to be taken as entries of a four-dimensional time series) are apparently unstructured or noisy. Can we find their underlying, hidden structure?

Fortunately, this problem of unidentifiability can be solved by independent component analysis (ICA) which assumes independent latent variables that are non-Gaussian. For example, applying ICA on the signal in Fig. 1, we “separate” them into four original source signals which were seriously mixed in the observed data, as shown in the Figure. This kind of application of ICA is often called “blind source separation.” If there were a factor rotation that remains undetermined, no such separation would be possible.

ICA will be reviewed in Sect. 3 and forms the basis of most of the theory reviewed in this paper.

1.2 Nonlinear representation learning and disentanglement

Machine learning has recently been profoundly transformed by deep learning, which essentially means learning (estimating) arbitrary nonparametric nonlinear functions from data. While the theory is quite well developed in the supervised case, consisting essentially of regression, the unsupervised case is much less developed. Unsupervised means that we only observe a potentially high-dimensional random vector \mathbf{x} , and there is no “output” or “label” or “regressand” defined (just like in factor analysis above). It is in fact widely appreciated that unsupervised, nonlinear learning of representations or features is one of the biggest challenges facing machine learning at the moment. It is often referred to as “disentanglement,” although this term is vague and not very well defined.

It is often assumed that most satisfactory solution for unsupervised deep learning is based on estimation of probabilistic generative models, because probabilistic theory often gives optimal objectives for learning, and immediately enables probabilistic inference of various quantities, such as the latent variables. This brings the machine learning theory in close connection with statistical estimation. Finding a good representation can then be defined as recovering the original latent variables that were assumed to generate the data. Some of the most popular probabilistic methods for unsupervised deep learning are based on variational autoencoders (VAE; Kingma and Welling 2014)

and generative adversarial networks (GAN; Goodfellow et al. 2014). These methods have been successful in approximating the probability density function (pdf) of the data and generating new data points.

Unfortunately, in most models used in unsupervised deep learning, there has not been any proof that the original latent variables can be recovered. In fact, they tend to rely on transformations of latent variables that are Gaussian and even white. We thus recover the same serious problem of unidentifiability (factor rotation) as in the case of linear models already mentioned.

Inspired by the theory of linear ICA, we might try to directly extend it to the nonlinear case. Such a basic framework would assume that the observed data is generated by an invertible nonlinear (nonparametric) transformation of non-Gaussian independent components, and we wish to recover them from observed data alone.

However, the extension of ICA to general nonlinear mixtures has proven very problematic. In particular, if the observed data \mathbf{x} are obtained as i.i.d. samples, i.e., there is no temporal or similar structure in the data, the model is seriously unidentifiable (Hyvärinen and Pajunen 1999). This is due to a further kind of unidentifiability that is specific to nonlinear models.

Fortunately, a solution to non-identifiability in nonlinear ICA can be found by utilizing temporal structure in time series (Harmeling et al. 2003; Sprekeler et al. 2014; Hyvärinen and Morioka 2016, 2017) or similar “auxiliary” information (Hyvärinen et al. 2019; Khemakhem et al. 2020a). Section 5 reviews the identifiability theory for such nonlinear latent variable model, in particular nonlinear ICA.

1.3 Causal discovery and structural equation models

Causal discovery is another goal of statistical analysis. It may appear to be unrelated to latent variable models, but we will see below that there is actually a deep connection. Causal models play a fundamental role in modern scientific endeavor (Spirtes et al. 2000; Pearl 2009; Peters et al. 2017). While randomized control studies are the gold standard to study the effect of one variable on another, such an approach is unfeasible or unethical in many scenarios (Spirtes and Zhang 2016). Furthermore, big data sets publicly available on the internet often try to be generic and thus cannot be strongly based on specific interventions. As such, it is both necessary and important to develop *causal discovery* methods through which to uncover causal structure from (potentially large-scale) passively observed data. Data collected without the explicit manipulation of certain variables is often termed *observational data*, in contrast to experimental data where certain variables are intervened upon, as in randomized controlled trials.

Before going to the general case, it is very useful to consider the linear bivariate case to understand the basic idea. The problem of causal discovery then essentially means finding the “direction of effect” by choosing between two regression models; either

$$x_2 = bx_1 + e_2 \quad (2)$$

which we might denote by $x_1 \rightarrow x_2$, or

$$x_1 = bx_2 + e_1 \quad (3)$$

which we might denote by $x_2 \rightarrow x_1$. It is well known that if the variables are standardized *Gaussian*, the situation is completely symmetric: The likelihood is equal for both models, and the variance explained is equal for both models (and so is the regression coefficient as already imposed above). Thus, we see again that for Gaussian variables we have a problem: the direction of effect is not identifiable. However, as will be seen below, it is identifiable for non-Gaussian variables.

In the general case, we use the framework of structural equation models² (SEMs) (Bollen 1989). Fundamentally, SEMs define a statistical model that describes the interactions of a set of observed variables $\mathbf{x} = (x_1, \dots, x_n)$ using a set of mutually independent disturbances or noise variables $\mathbf{e} = (e_1, \dots, e_n)$. However, SEM is not only a statistical model of a probability distribution, but also a mathematical tool that can be used to encapsulate causal knowledge (Pearl 2009). SEMs are in this sense more powerful than latent variable models: not only do they describe the set of all distributions, they can be used to perform interventions and answer counterfactual queries by changing the noise distribution or the causal mechanism in one or more of Eq. 23). We leave a formal definition for later, and just illustrate the SEM by Fig. 2 which shows how the influences of the observed variables can be described by a graph, where the arcs show which x_i causally influences which x_j .

SEM are only useful for causal discovery if they define an *identifiable* causal model. In the case of a causal model, identifiability typically means that we can distinguish between cause and effect, or find the direction of effect as in the example above. In general, it means we can find the right “causal ordering” over the variables (x_1, \dots, x_n) such that, loosely speaking, the earlier variables are the causes of the later variables. Without identifiability, different causal orders would give rise to the same data, which would prevent any real scientific analysis of causality.

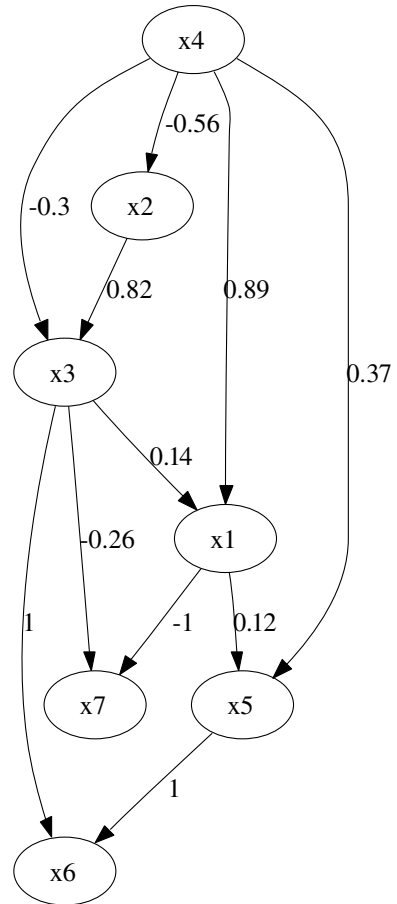
Although the problem of causal discovery may seem, at first sight, very different from estimating latent variable models, the deep result here is that we can in some cases reduce the estimation of a SEM to the estimation of a latent variable model such as ICA. Thus, if we develop an identifiable latent variable model, we may be able to develop a corresponding SEM that is also identifiable. In fact, even the estimation methods can be largely shared. Section 4 will consider identifiability in the linear case, while Sect. 6 will consider the nonlinear case.

2 Definition of identifiability

Now we shall proceed to a formal definition of identifiability. Identifiability is an important property of probabilistic models. For example, it seems futile to interpret the quantities estimated in the model, whether parameters or some further latent variables, if the model is not identifiable. We defer a more detailed discussion of the utility of identifiability to Sect. 7, and simply give the definition and some examples here.

² Structural equation models (SEM) are sometimes referred to as structural causal models (SCM) or functional causal models (FCM) in recent machine learning literature.

Fig. 2 A SEM can be expressed by a directed graph (typically acyclic), where the arcs express causal influences, as well as statistical dependencies. Here, the nodes have been ordered so that the influences all go from top to bottom. The disturbance or noise variables are not plotted here, since each observed variable simply has its own disturbance variable



Consider a probabilistic model \mathcal{P} , defined as the set of distributions $\{P_\theta : \theta \in \Theta\}$ with parameter θ taking values in some set Θ , on a set of possible observations \mathcal{X} (typically \mathbb{R}^n in our case). A model \mathcal{P} is said to be *identifiable* if the mapping $\theta \in \Theta \mapsto P_\theta(\mathbf{x})$ is injective (i.e., one-to-one):

$$(P_{\theta_1}(\mathbf{x}) = P_{\theta_2}(\mathbf{x}), \forall \mathbf{x} \in \mathcal{X}) \implies \theta_1 = \theta_2. \quad (4)$$

In other words, if two parameters θ_1 and θ_2 generate the same distribution over the set of observations \mathcal{X} , then they are necessarily equal. It is important to note that identifiability is a property of the probabilistic model and not of any particular estimation method.

As a simple theoretical example, consider that we flipped a possibly biased coin N times. A coin toss only has two outcomes: heads with probability $\theta \in [0, 1]$, or tails with probability $1 - \theta$. Let us define two different models for this data. In the first case, we use a Bernoulli model with parameter θ for the outcome of a coin flip; denote by $\text{Ber}_\theta(x)$ the probability of the coin flip given $x \in \{H, T\}$ which

codes heads and tails. This model is identifiable. To see this, let (θ_1, θ_2) be such that $\text{Ber}_{\theta_1}(x) = \text{Ber}_{\theta_2}(x)$ for $x \in \{H, T\}$. Since $\text{Ber}_{\theta}(H) = \theta$, and this probability is directly given by the data in the limit of infinite coin flips, we conclude that $\theta_1 = \theta_2$ and that the model is identifiable. On the other hand, as a counterexample, we can imagine something like a latent variable model: we do not observe a real coin, but rather the output of a computer simulation. Tossing a coin in this case proceeds in two steps: we first draw a sample $z \sim \mathcal{N}(\mu, \sigma^2)$, i.e., a latent variable with a Gaussian distribution with mean μ and variance σ^2 ; then we assign heads to x if $z \geq 0$ and tails otherwise. This is effectively a latent variable model with parameters (μ, σ) . Crucially, elementary calculations show that the observed probabilities depend only on the ratio $\frac{\mu}{\sigma}$, which takes the same value for infinitely many pairs (μ, σ) , which means that this model is not identifiable.

As a most fundamental example for our purposes, consider the *Gaussian factor analysis* in Eq. (1), where the components or factors s_i are Gaussian, uncorrelated, and have unit variance. This model is well known not to be identifiable. For simplicity, consider the matrix \mathbf{A} to be square and orthogonal— denote it by \mathbf{U} to highlight those assumptions—and ignore the noise \mathbf{n} . An intuitive justification for unidentifiability would be that a Gaussian distribution is completely determined by covariances (and means). Now, the number of covariances is $\approx n^2/2$ due to symmetry, so we cannot solve for the n^2 parameters in the mixing matrix as we have “more variables than equations.” More rigorously, the Gaussian distribution exhibits a rotational symmetry when the covariance matrix is identity (as is the case for the components here). The pdf is given by the probability transformation formula (transforming \mathbf{x} back to \mathbf{s}):

$$p(\mathbf{x}) = \frac{1}{(2\pi)^{n/2}} \exp\left(-\frac{1}{2}\|\mathbf{U}^T \mathbf{x}\|^2\right) |\det \mathbf{U}^T| = \frac{1}{(2\pi)^{n/2}} \exp\left(-\frac{1}{2}\|\mathbf{x}\|^2\right) \quad (5)$$

where the last equality comes from the orthogonality of \mathbf{U} . Now, we see that the pdf of \mathbf{x} does not depend on \mathbf{U} at all. Thus, \mathbf{U} cannot be identifiable. In practice, we could rotate the factors \mathbf{s} by any orthogonal matrix, and compensate for that by rotating the columns of \mathbf{U} by the inverse, and the data distribution would stay the same, which is why this is called the “factor rotation indeterminacy.”

The basic definition of identifiability talks thus about identifiability of parameters. Sometimes, by slight abuse of terminology, we also talk about *identifiability of the latent variables*, but it is not quite clear how that should be defined. If we consider the factor analysis model in Eq. (1) without noise, knowing the matrix \mathbf{A} will immediately give us the components \mathbf{s} by using the (pseudo)inverse of \mathbf{A} (assuming $m \geq n$ as is typical). Thus, identifiability of \mathbf{A} implies, in a non-rigorous sense, identifiability of \mathbf{s} . On the other hand, if there actually is noise in the model, knowing \mathbf{A} will not give us the components \mathbf{s} since the noise cannot be completely removed. In this case, identifiability could be defined in the sense that the posterior $p(\mathbf{s}|\mathbf{x}, \mathbf{A})$ can be recovered in a suitable sense (Khemakhem et al. 2020a). However, we leave such a definition aside here, and focus on identifiability of the parameters.

It should also be noted that the strict definition of identifiability in Eq. (4) may be limiting in some cases. In practical scenarios, we may want to introduce

identifiability with slightly relaxed definitions, such as identifiability of parameters up to an equivalence class. For example, independent components are only identifiable up to arbitrary scaling of the components, but this is usually not considered a problem.

3 Linear independent component analysis

ICA is a statistical latent variable model which is very closely related to the classic factor analysis model reviewed above. The basic idea is that assuming the components to be non-Gaussian breaks the rotational symmetry just described and leads to an identifiable model. In particular, for non-Gaussian data, higher-order moments give more information than that contained in the covariances (Hyvärinen and Oja 2000; Hyvärinen et al. 2001). The identifiability of the basic ICA model is well known since Comon (1994), and in fact was proved in the 1950's by Darmois and Skitovich.

3.1 Definition and identifiability

The basic model is as follows (Comon 1994; Jutten and Héroult 1991; Hyvärinen and Oja 2000; Hyvärinen et al. 2001). Assume $s_i, i = 1, \dots, n$ are n independent, non-Gaussian latent random variables. Assume \mathbf{A} is an invertible $n \times n$ matrix. We observe the random vector \mathbf{x} which is a linear mixture given by

$$\mathbf{x} = \mathbf{A}\mathbf{s}. \quad (6)$$

We assume that the means of the s_i are centered to zero and that their variances are finite and normalized to unity. Note the differences to the factor analysis model:

- (1) The components are assumed to be *non-Gaussian*. This is the fundamental difference which sets the model apart from classical factor analysis. (While the first letter in “ICA” emphasizes independence, the factors in classical Gaussian factor analysis are independent as well; the terminology is slightly misleading regarding this difference).
- (2) The number of components (factors) is equal to the number of observed variables. This is basically assumed for the sake of mathematical simplicity: the matrix \mathbf{A} is then uniquely invertible, but the assumption could be relaxed. At the same time, it emphasizes the fact that in the context of ICA we are not so interested in dimension reduction but on finding the original components. In practice, the dimension is often reduced by principal component analysis (PCA) before application of ICA. In terms of the classical terminology of factor analysis, ICA

can then be seen as a “factor rotation,” considering the principal components as estimates of factors.

3. There is no noise term. This is partly justified by the large number of components: some of them can be noise, whereas others are more interesting components. Since we do not reduce the dimension of the data, all variance of the data is “explained” by the components anyway.

It is really the non-Gaussianity that fundamentally distinguishes the ICA model from classical factor analysis, and enables identifiability. The importance of non-Gaussianity is further emphasized by the theory of ICA estimation, which provides also an intuitive proof of identifiability as follows. One of the most fundamental theorems in ICA says that we can estimate ICA by finding an invertible transformation that maximizes the non-Gaussianity of the components. In fact, each component corresponds to a maximum (over some parameters w_i) of non-Gaussianity of a linear combination $\sum_i w_i x_i$. This is because, loosely speaking, by the central limit theorem, a sum of independent random variables is more Gaussian than any of the original random variables (strictly speaking this holds if the random variables have the same distribution). The linear combination $\sum_i w_i x_i$ is a linear combination of the s_i , and thus its non-Gaussianity is maximized when it is actually equal to one of the s_i . See Hyvärinen and Oja (2000), Hyvärinen et al. (2001) for details.

Note two rather trivial indeterminacies, i.e., unidentifiable aspects of the ICA model. First, the ordering of the components is not identifiable, or even defined by the model. Second, each component can only be estimated up to linear scaling (and sign), since if a component is multiplied by a scalar constant and the corresponding column of \mathbf{A} is divided by that constant, the data distribution stays the same. This indeterminacy is partly removed when the variances of the components are conventionally defined to be equal to unity, which also means that \mathbf{s} is white, but this is a mere convention and the real variance of the components is unidentifiable.

3.2 A simple identifiability proof

Next, we provide a simple rigorous identifiability proof for ICA. The price to pay for its simplicity is some reduction in generality compared to the celebrated Darmois–Skitovich theory. In our simple case, the identifiability theorem takes then the following form:

Theorem 1 *Assume the independent components have finite variance, and the log-pdf's of the independent components have continuous second derivatives. If the variables $x_i, i = 1, \dots, n$ in Eq. (6) are mutually independent, \mathbf{A} has exactly one nonzero entry in each row and each column.*

Proof ³ As typical in ICA theory, we can assume without restriction of generality that \mathbf{A} is orthogonal. This is because \mathbf{s} is white and we also can whiten \mathbf{x} as preprocessing, which implies the orthogonality of \mathbf{A} since the covariance of \mathbf{x} is equal to $\mathbf{A}\mathbf{A}^T$.

The pdf of \mathbf{x} is obtained by the formula for transformation of probability density as

$$p(\mathbf{x}) = \prod_{i=1}^n p_i \left(\sum_{j=1}^n a_{ji} x_j \right) |\det \mathbf{A}^T| \quad (7)$$

where p_i denotes the pdf of the i -th component. Taking logarithms, and noting that the determinant of an orthogonal matrix is equal to ± 1 , we get

$$\log p(\mathbf{x}) = \sum_{i=1}^n \log p_i \left(\sum_{j=1}^n a_{ji} x_j \right). \quad (8)$$

Next, since the x_i are assumed mutually independent, $\log p(\mathbf{x}) = \sum_{i=1}^n f_i(x_i)$ for some functions f_i . This implies that any second-order cross-derivative must be zero:

$$\frac{\partial^2 \log p(\mathbf{x})}{\partial x_k \partial x_l} = \sum_{i=1}^n a_{ki} a_{li} (\log p_i)'' \left(\sum_{j=1}^n a_{ji} x_j \right) = 0, \quad \text{for all } k \neq l. \quad (9)$$

This set of equations can be collected together and expressed in matrix form as

$$\mathbf{A}^T \text{diag}_i \left[(\log p_i)'' \left(\sum_{j=1}^n a_{ji} x_j \right) \right] \mathbf{A} = \text{diag}_i [c_i(\mathbf{x}; \mathbf{A})] \quad (10)$$

where the $c_i(\mathbf{x}; \mathbf{A})$ are some unknown scalar-valued functions of \mathbf{x} and \mathbf{A} corresponding to the case $k = l$. This equation must hold for all \mathbf{x} .

For Gaussian densities (and no others), $\log p_i$ is quadratic, and its second derivative is constant. We have assumed non-Gaussianity, so $(\log p_i)''$ is not constant. By continuity, $(\log p_i)''$ takes in fact an infinity of different values, since it takes all the values in some interval of the real line. This implies we can find for each i a point y_i such that the entries $d_i = (\log p_i)''(y_i)$ are all distinct, i.e., $d_j \neq d_k$ for $j \neq k$. (This is possible even if one of the components is Gaussian.) By the invertibility of \mathbf{A} , we can find corresponding \mathbf{x} such that $y_i = \sum_{j=1}^n a_{ji} x_j$ for all i . In the following, we consider (10) with \mathbf{x} fixed to such value, so the diagonal entries d_i and c_i are fixed.

³ Alternatively, we could do the same proof in the Fourier domain, i.e., using characteristic functions \hat{p} . Then, $p(\mathbf{x})$ and $p_i(s_i)$ will be replaced by the characteristic functions, and the Jacobian disappears in the first equations. Thus, we would replace the assumption of smooth pdf's by the assumption of continuous second derivatives of the characteristic functions of the p_i , denoted by \hat{p}_i . Such an assumption is related to the moment structure of the components: it is just slightly more restrictive than assuming finite variances for the components. The whole proof is valid for characteristic functions with minimal changes. Thus, we get a more general if a bit more complicated proof.

Now, the equation on the LHS of (10) is actually an eigenvalue decomposition (EVD), since \mathbf{A} is orthogonal. Importantly, the diagonal entries (eigenvalues) are distinct, which implies by a well-known results in linear algebra that the EVD is unique up to ordering of the eigenvalues and the signs of the eigenvectors. The RHS can be interpreted as an EVD as well, with eigenvectors such that each row and column has exactly one nonzero entry. Since the “eigenvectors” of the EVD on both sides must match up to permutation, the eigenvectors of \mathbf{A} must have the same property for nonzero entries. \square

Note that if the components were Gaussian, the diagonal matrix on the left-hand side of (10) would be equal to identity, and thus the whole LHS is equal to identity for any \mathbf{A} . Then, any orthogonal \mathbf{A} fulfills (10), since it is easy to show that the right-hand side must necessarily be identity as well. This shows how Gaussian variables are not allowed. In fact, the uniqueness of the eigenvalue decomposition only holds for distinct eigenvalues.

Another point to note is that the matrix \mathbf{A} is not shown to be identity. The form that \mathbf{A} takes enables a permutation of the components, as well as multiplying each by a scalar constant. These are fundamental indeterminacies that cannot be solved in any model similar to factor analysis, as already pointed out above.

3.3 Alternative approaches

The literature of blind source separation actually considers non-Gaussianity as only one, if the most important, principle that enables estimation of a linear mixing model. In the case of time series data, identifiability is enabled by two very different kinds statistical structure: autocorrelations of the components (Tong et al. 1991; Belouchrani et al. 1997) or their nonstationarity (Matsuoka et al. 1995; Pham and Cardoso 2001). Together with non-Gaussianity, these constitute what Cardoso (2001) called the “three easy routes to [linear] ICA.” We shall consider these principles in more details in the case of nonlinear ICA below.

Yet another approach is possible in the case where the data is nonnegative, and in particular so that the data has a concentration to small positive values. This leads to the model originally called positive matrix factorization (Paatero and Tapper 1994) and popularized under the heading of nonnegative matrix factorization (NMF) by Lee and Seung (1999). Its identifiability has been analyzed by Donoho and Stodden (2004). Combinations of NMF and ICA are considered by Plumbley (2003), Hoyer (2004). Related to this, Hyttinen et al. (2022) considered the case where the observations are binary.

We further note that great similarity of the ICA model with dictionary learning and sparse coding. The similarity was well known in early work of those methods (Olshausen and Field 1997) but seems to have been forgotten recently. The only real differences between ICA and a probabilistic formulation of dictionary learning are that in dictionary learning, the number of components is large ($n > m$), and there is noise. Some identifiability results for such cases are considered by Eriksson and Koivunen (2004).

4 Linear structural equation model

Next, we consider a linear SEM, which we use here as a fundamental framework for causal discovery. We will see how its identifiability can be proven based on the theory of linear ICA.

4.1 Definition of model

A linear SEM consists of a collection of equations of the form

$$x_i = \sum_{j=1}^n b_{ij}x_j + e_i, \quad i = 1, \dots, n, \quad (11)$$

where the x_i are the observed variables, and the e_i are latent variables called the disturbances, external influences, or simply noise variables. The idea is that the variable x_i is caused by those x_j for which the coefficient b_{ij} is nonzero.

An SEM is often associated with a directed acyclic graph (DAG) \mathcal{G} called the *causal graph*. Each node of \mathcal{G} corresponds to an observed variable x_i , and there is an edge from x_j to x_i iff b_{ij} is nonzero. We have here actually made the assumption of acyclicity very typical in this context. It means that when the matrix \mathbf{B} is interpreted as a connection matrix of a (weighted) graph, the ensuing graph has no cycles: one cannot start at a node and follow the edges so that one comes back to the starting node. An interesting property of a DAG is that the nodes (variables) can be ordered so that all the connections go “forward” in that ordering; this is called a “causal order,” but it is not necessarily unique. An illustration of such a DAG expressing a SEM, visually ordered according to the causal ordering, was given in Fig. 2.

The problem is now to estimate the parameters b_{ij} based on observations of \mathbf{x} . It is well known that the problem is in general ill-posed. In particular, for Gaussian data the model is unidentifiable, as shown above for the case of two variables. One wide-spread solution is to use prior knowledge to fix most of the b_{ij} to zero; for example, randomized controlled trials might be available to provide that knowledge. However, using such prior knowledge, not to mention interventions, is in strong contradiction with the aim of *discovery* which is central to us.

4.2 Identifiability

The key to identifiability of the SEM is its close relation to ICA. This intimate relation of latent variable models and SEMs is a central theme of this paper, and we will later see how it applies even in the nonlinear case.

Denote by \mathbf{B} a matrix which collects all the coefficients b_{ij} . We can express the SEM as

$$\mathbf{x} = \mathbf{B}\mathbf{x} + \mathbf{e} \quad (12)$$

where the vector \mathbf{e} collects the external influences. Now, by elementary linear algebra this implies

$$\mathbf{x} = (\mathbf{I} - \mathbf{B})^{-1} \mathbf{e}. \tag{13}$$

In other words, the SEM implies that the data follows a latent variable model. In fact, this is nothing else than an ICA model with mixing matrix $\mathbf{A} = (\mathbf{I} - \mathbf{B})^{-1}$, under suitable assumptions. In particular, assume that the e_i are *mutually independent and non-Gaussian*. Then, the model in Eq. (13) is exactly an ICA model, and has by definition the same number of components (corresponding to disturbances) as observed variables.

Thus, it would seem that we can estimate an ICA model of \mathbf{x} , and transform the obtained \mathbf{A} back to \mathbf{B} by simply $\mathbf{B} = \mathbf{I} - \mathbf{A}^{-1}$. However, there is one serious complication: ICA does not estimate order of the components, e_i . As pointed out in Sect. 3, there is no ordering inherently defined between the components. This is in stark contrast to the SEM, where we know that e_i is the external influence of the variable x_i (for the same index i). Thus, we need to find the right ordering of the independent components, which implies an ordering of the columns of the mixing matrix \mathbf{A} , before we can transform back to \mathbf{B} .

The recovery of the correct ordering is not possible in general. However, as mentioned above, in the theory of SEM and causal discovery the assumption of a directed *acyclic* graph (DAG) is often made. Based on the acyclicity assumption, the right ordering of the components can be found. To see how this is possible, consider the 2D case. Assume we have estimated the inverse of the mixing matrix $\mathbf{W} = \mathbf{A}^{-1}$, for some arbitrary ordering of the components. Based on Eq. (13), we see that we should have $\mathbf{W} = \mathbf{I} - \mathbf{B}$. The acyclicity of \mathbf{B} means in this 2D case that it has exactly one nonzero entry, by definition in the off-diagonal (unless the graph is degenerate and \mathbf{B} is all zeros). Thus, \mathbf{W} has exactly one zero entry. Denoting an arbitrary nonzero entry by $*$, the real \mathbf{W} could, for example, be of the form

$$\mathbf{W} = \begin{pmatrix} 1 & * \\ 0 & 1 \end{pmatrix}. \tag{14}$$

Now, if the rows of \mathbf{W} are switched to the wrong order in the estimation due to the indeterminacy of the order of components, it is easy to see that the zero goes to the diagonal. This is a contradiction since the diagonal entries are all one by definition (even allowing for any rescaling). Thus, among the two different orderings, it is possible to find the right one based on acyclicity. The general case is explained by Shimizu et al. (2006); note that acyclicity is sufficient but not necessary (Lacerda et al. 2008).

A smaller detail is that the normalization of the mixing matrix in ICA and SEM is different: In ICA, the variances of the components are typically defined to be unity, while in SEM, a related normalization is obtained by the fact that $\mathbf{I} - \mathbf{B}$ has all ones in the diagonal. However, this is just a normalization convention that has little implication for identifiability.

Thus, we see that the SEM is identifiable under the assumptions of independence and non-Gaussianity (like in ICA), together with the new assumption of acyclicity. The

resulting model is called LiNGAM for linear non-Gaussian acyclic model. We refer the reader to Shimizu et al. (2006) for details on this basic model, and Shimizu (2014) for a longer treatment with some more recent developments.

Regarding estimation of LiNGAM, it is also possible to develop a very explicit interpretation of the resulting SEM estimation in terms of maximization of non-Gaussianity (Hyvärinen and Smith 2013), just like in the case of ICA estimation.

4.3 Alternative approaches

Alternative identifiable SEM frameworks have been proposed, e.g., by (Peters and Bühlmann 2014; Jakobsen et al. 2022). These remove the non-Gaussianity assumption but introduce alternative restrictions more inspired by causal inference literature. One particularly important problem with causal discovery is that we may not observe all the relevant variables; discovery in the presence of such hidden “confounders” is considered, e.g., by Hoyer et al. (2008), Tashiro et al. (2014). A combination of source separation and LiNGAM was proposed by Monti and Hyvärinen (2018). Finally, we note that the alternative methods for estimating a linear mixing model, reviewed in Sect. 3.3 above, could also be used to estimate a linear SEM. In fact, Zhang and Hyvärinen (2010) use a related idea to develop another combination of source separation and causal discovery.

5 Nonlinear independent component analysis

Nonlinear ICA is a fundamental problem in unsupervised learning which has attracted a considerable amount of attention recently. It promises a principled approach to representation learning and “disentanglement,” in particular using deep neural networks. Nonlinear ICA attempts to find nonlinear components in multidimensional data by generalizing the linear ICA framework in Sect. 3. The essential difference to most methods for unsupervised representation learning is that the approach starts by defining a generative model in which the original latent variables can be recovered, i.e., the model is identifiable by design.

5.1 Problems in identifiability

We start by introducing a simple model for nonlinear ICA which turns out to be unidentifiable. Denote, as above, an observed n -dimensional random vector by $\mathbf{x} = (x_1, \dots, x_n)$. We assume it is generated using n independent latent variables called independent components, s_i . A straightforward definition of the nonlinear ICA problem is to assume that the observed data is an arbitrary (but smooth and invertible) transformation \mathbf{f} of the latent variables $\mathbf{s} = (s_1, \dots, s_n)$ as

$$\mathbf{x} = \mathbf{f}(\mathbf{s}). \quad (15)$$

The goal is then to recover the inverse function \mathbf{f}^{-1} as well as the independent components s_i based on observations of \mathbf{x} alone.

Research in nonlinear ICA has been hampered by the fact that such simple approaches to nonlinear ICA are not identifiable, in stark contrast to the linear ICA case. To put it simply, for any x_1, x_2 , one can always find a function $g(x_1, x_2)$ which is independent of x_1 . Darrow provided such a construction back in 1952, thus showing the “impossibility” of nonlinear ICA. He simply used the conditional cumulative distribution function

$$g(\xi_1, \xi_2) = P(x_2 < \xi_2 | x_1 = \xi_1) \quad (16)$$

to construct a new variable $z = g(x_1, x_2)$ which turns out to be independent of x_1 . (A slight generalization was provided by Hyvärinen and Pajunen (1999), and a slight variation by Locatello et al. (2019).) The problem is that we can apply this equally well after making an initial transformation $\tilde{x}_1 = h(x_1, x_2)$, so we see that *any* such transformation \tilde{x}_1 could be considered an independent component, since we can always find a decomposition of the data as \tilde{x}_1 and the variable defined by the construction above. The marginal distributions of the two variables can further be easily transformed to anything we like, so we can, in particular, reproduce the *distributions* of the original independent components exactly without recovering the *components* themselves. (This logic is a bit different from the definition of identifiability given above, since here we consider the distributions of the latent variables, but it is equivalent.) Thus, it is clear that independence is not a strong enough assumption to enable identifiability, whatever assumptions (e.g., non-Gaussianity) we may make on the distributions of the independent components s_i .

The unidentifiability is of course a major problem in practice since most of the utility of ICA rests on the fact that the model is identifiable, or in alternative terminology, the “sources can be separated.” We should also note that curiously, many unsupervised deep learning models, such as VAE and GAN, assume a latent variable model where the s_i are even Gaussian, thus recreating even the classical Gaussian factor rotation problem and making the unidentifiability even worse.

5.2 Identifiable model definition

However, it should be noted that Darrow’s counterexample above assumes that the data is sampled i.i.d. (independently and with identical distributions) from a random vector; in particular, the observations (data points) are statistically independent of each other (which is not to be confused with the independence of the components). A promising direction is to relax the assumption of i.i.d. sampling. A fundamental case is to consider time series, and the *temporal structure* of independent components. Thus, we assume

$$\mathbf{x}(t) = \mathbf{f}(\mathbf{s}(t)) \quad (17)$$

where t is the time index. As a first attempt, we can assume that the sources $s_i(t)$ have nonzero *autocorrelations*, which has a long history in the linear case (Tong et al. 1991; Belouchrani et al. 1997). In machine learning literature, such models have been widely used under the heading of “temporal coherence,” “temporal stability,” or “slowness” of the features. Harmeling et al. (2003) proposed that we could

try to find nonlinear transformations which are maximally uncorrelated even over time lags and after nonlinear scalar transformations. Sprekeler et al. (2014) showed that a closely related method enables separation of sources if they all have distinct autocorrelations functions. Sprekeler et al. (2014) thus constituted probably the first identifiability proof for nonlinear ICA with general nonlinearities; however, it suffers from the restrictive condition that the sources must have different statistical properties, which is rather unrealistic in many cases. Note that just like the linear scaling is unidentifiable in linear ICA, a nonlinear scaling by a monotonic function is fundamentally unidentifiable in nonlinear ICA; in both cases, the global sign of each component and their ordering are unidentifiable as well.

Hyvärinen and Morioka (2017) proposed a rigorous and general treatment of identifiability for components which are stationary time series, the components being independent from each other, but each component having *temporal dependencies*. Assuming the components $s_i(t)$ have sufficient temporal dependencies and sufficient non-Gaussianity according to certain technical definitions, the model was proven to be identifiable, up to monotonic pointwise transformations of the components. The proof was refined by Hälvä et al. (2021), partly based on Schell and Oberhauser (2023); an alternative approach was proposed by Klindt et al. (2020).

Intuitively speaking, temporal dependencies can provide much more information since now the components are independent over any lags, i.e., $s_i(t)$ and $s_j(t - \tau)$ are independent for any lag τ . Such independence can then be imposed on any estimates of the components. This provides many more constraints compare to the i.i.d. case considered in the Darmois construction given above (Schell and Oberhauser 2023). Therefore, it is intuitively plausible that the model becomes identifiable.

Another form of temporal structure that has been previously used in the case of linear blind source separation is *nonstationarity* (Matsuoka et al. 1995, Pham and Cardoso 2001), in particular nonstationarity of variance. Such nonstationarity of variances seems to be prominent in many kinds of real data, for example, in EEG and MEG (Brookes et al. 2011), natural video (Hyvärinen et al. 2009), and closely related to changes in volatility in financial time series. This principle was extended to the nonlinear case by Hyvärinen and Morioka (2016), Khemakhem et al. (2020a) gave the strongest identifiability results based on this principle so far. They assumed a piecewise nonstationary model where the components in each segment follow an exponential family, while the parameters of the exponential family change from one segment to another. It was further assumed that the segmentation is known. Then, the identifiability of the model was proven under technical assumptions that guarantee the nonstationarity is strong enough. While the assumptions here look quite restrictive, this result can be extended in various ways as will be seen below. Note that in this theory, the components can usually only be recovered up to a nonlinear, pointwise but not necessarily monotonic, function of the components; for example, the only squares of the components can be recovered in the simplest models.

Again, we can give the identifiability results a simple intuitive justification. By the basic assumption of independence, the components are independent at any time point, i.e., $s_i(t)$ and $s_j(t)$ must be independent for any t . Crucially, nonstationarity implies that the distributions inside the segments are different; thus, the number of independence constraints is basically multiplied by the number of segments. This

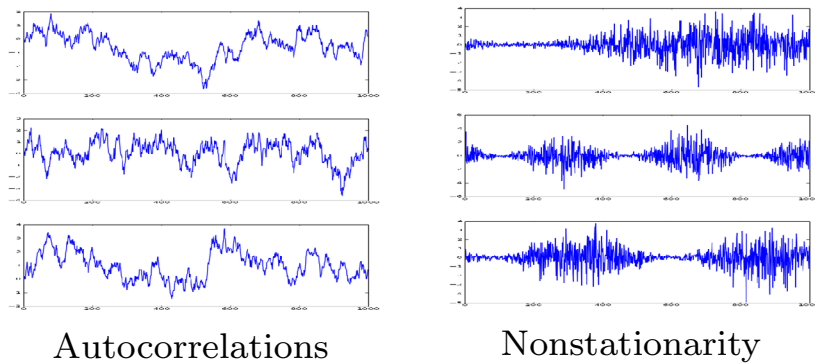


Fig. 3 The two fundamental kinds of temporal structure used in nonlinear ICA. Left: Autocorrelated sources, which show the basic form of temporal dependencies. Right: Nonstationary sources, in particular of exhibiting nonstationarity of variances. If the independent components have either of such structures, the model can be identifiable. Furthermore, such structures can be combined or generalized in various ways

way we get many more constraints compared to the i.i.d. case, which was considered in the Darmois construction given above.

The two kinds of temporal structure used here are illustrated in their basic forms in Fig. 3. Such a dichotomy of temporal structure is useful for intuition about the temporal dependencies. However, nonstationarity can be subsumed as a case of temporal dependencies by assuming that the apparent nonstationarity stems from a hidden Markov model (Hälvä and Hyvärinen 2020). This leads to the general theory of Hälvä et al. (2021) which shows identifiability for very general temporal dependencies, including those coming from HMM's. A related model combining the two kinds of temporal structure was proposed by Morioka et al. (2021).

Another way of generalizing identifiability is to realize that the proofs based on nonstationarity simply assume that conditionally on the segment index, the distributions of the components change, while they impose no temporal structure inside the segments. Thus, we can replace this segment index by any observed *auxiliary variable* that conditions the distributions of the components. It could be the index of a subject in a biomedical setting, the index of a country in an international study, or anything that changes the underlying statistics in a suitable way. Thus, we get another set of identifiability results, based on observation of such an additional auxiliary information (Hyvärinen et al. 2019). Khemakhem et al. (2020a, 2020b) generalize those results even further; see also Zimmermann et al. (2021). We will see below how this is particularly useful in the case of causal discovery.

Another generalization of the results is to consider the noisy mixing model, where noise is added to the mixing as in the classical factor analysis in Eq. 1. Such additive noise has some algorithmic advantages and has actually been widely used in related unsupervised models as well as by Khemakhem et al. (2020a). While its distribution is typically considered known, Hälvä et al. (2021) proved the identifiability of the noise distribution under quite general conditions.

5.3 Sketch of an identifiability proof

Next we provide a sketch of a simple identifiability proof following Hyvärinen and Morioka (2016). This is only a special case of the currently available theory, but serves to illustrate the basic principles. For the most sophisticated currently available proofs based on Hyvärinen and Morioka (2016), the reader is referred to Khemakhem et al. (2020a, 2020b), whose proofs are thus generalizations of the following. Note also that Hälvä et al. (2021) developed a proof using a very different method, based on Hyvärinen and Morioka (2017), and reminiscent of the linear ICA identifiability proof given above. Another very different approach to the proofs was proposed by Hyvärinen et al. (2019).

We consider a nonstationary model where at each segment $\tau = 0, \dots, T$, the i -th component follows an exponential family of order one:

$$\log p_{\tau}(s_i) = b_i(s_i) + \lambda_{\tau,i} q_i(s_i) \quad (18)$$

with the q_i, b_i being the sufficient statistic and the base measure of component i , and $\lambda_{\tau,i}$ the parameter for component i and segment τ . (Instead of the segment τ , the proof could also be developed for a time index t with no changes.) Now, compute the log-pdf of a data point \mathbf{x} in the segment τ under the nonlinear ICA model. Using the probability transformation formula, the log-pdf is given by

$$\log p_{\tau}(\mathbf{x}) = \sum_{i=1}^n b_i(g_i(\mathbf{x})) + \lambda_{\tau,i} q_i(g_i(\mathbf{x})) + \log |\det \mathbf{J}\mathbf{g}(\mathbf{x})| \quad (19)$$

where we drop the index t from \mathbf{x} for simplicity, and $\mathbf{g}(\mathbf{x}) = (g_1(\mathbf{x}), \dots, g_n(\mathbf{x}))^T$ is the inverse function of (the true) mixing function \mathbf{f} ; thus, $s_i = g_i(\mathbf{x})$ by definition. \mathbf{J} denotes the Jacobian matrix.

Assume we now have two different models which give the same data distribution, thus violating identifiability. That is, there is another inverse function $\tilde{\mathbf{g}}$, sufficient statistics and base measure \tilde{q}, \tilde{b} and parameters $\tilde{\lambda}$ such that for all τ ,

$$\begin{aligned} \sum_{i=1}^n b_i(g_i(\mathbf{x})) + \lambda_{\tau,i} q_i(g_i(\mathbf{x})) + \log |\det \mathbf{J}\mathbf{g}(\mathbf{x})| \\ = \sum_{i=1}^n \tilde{b}_i(g_i(\mathbf{x})) + \tilde{\lambda}_{\tau,i} \tilde{q}_i(\tilde{g}_i(\mathbf{x})) + \log |\det \mathbf{J}\tilde{\mathbf{g}}(\mathbf{x})|. \end{aligned} \quad (20)$$

Now, subtract both sides of this equation for the corresponding terms obtained for $\tau = 0$. We get

$$\sum_{i=1}^n \alpha_{\tau,i} q_i(g_i(\mathbf{x})) = \sum_{i=1}^n \tilde{\alpha}_{\tau,i} \tilde{q}_i(\tilde{g}_i(\mathbf{x})) \quad (21)$$

with $\alpha_{\tau,i} = \lambda_{\tau,i} - \lambda_{0,i}$. Remarkably, we got rid of the Jacobian terms.

Now it is enough to collect Eq. (21) for all $\tau = 1, \dots, T$. Assume that the matrix of the α is full rank, which means that we have enough segments (their number is

greater than the dimension of the data) and there is enough variability in their distributions. Then, we can solve for the $\tilde{q}_i(\tilde{g}_i(\mathbf{x}))$ which are then given by a linear transformation of the $q_i(g_i(\mathbf{x}))$. Such a linear indeterminacy can be resolved by linear ICA assuming that the components are marginally independent (in addition to conditionally independent in each segment). Thus, we see that $\tilde{q}_i(\tilde{g}_i(\mathbf{x})) = q_j(g_j(\mathbf{x})) = q_j(s_j)$ for some permutation of the indices i, j , and we obtain the components up to the pointwise nonlinear transformation given by the sufficient statistics. (The pointwise transformation is characterized in detail by Khemakhem et al. (2020a), who actually do not need a condition of marginal independence.) \square

5.4 Alternative approach by constraining nonlinearity

An alternative approach to making nonlinear ICA identifiable consists of constraining the mixing function. A natural approach would be to assume that the mixing function is close to linear (Zhang and Chan 2008). Recent research has focused on imposing constraints on the Jacobian matrix \mathbf{Jf} of the mixing function to achieve something to that effect.

The classic Liouville's theorem of conformal mappings is highly relevant here. It considers the case where the Jacobian is almost orthogonal in the sense that

$$\mathbf{Jf}(\mathbf{x})^T \mathbf{Jf}(\mathbf{x}) = \alpha(\mathbf{x})\mathbf{I} \quad (22)$$

for a scalar-valued function α . The theorem states that any such sufficiently smooth function \mathbf{f} has to belong to a very restricted class of functions (called Möbius transformations). A simple corollary of the theorem, stated and proven in Appendix A, says that if $\alpha \equiv 1$, i.e., the Jacobian is orthogonal, the function \mathbf{f} is actually necessarily affine: $\mathbf{f}(\mathbf{x}) = \mathbf{U}\mathbf{x} + \mathbf{b}$ for some orthogonal matrix \mathbf{U} . This theory thus provides a strong result on what kind of constraints on the Jacobian are meaningful. Constraining the Jacobian to be orthogonal is not meaningful since it does not allow for any nonlinear mixing functions. An intriguing question is, therefore, to explore some relaxations of the constraint in Eq. (22). Hopefully, some such relaxations would allow for a sufficiently large class of nonlinear functions, while still providing identifiability (Gresele et al. 2021; Zimmermann et al. 2021; Buchholz et al. 2022).

A number of further approaches constraining the mixing have recently been proposed. Kivva et al. (2022) showed identifiability in the case of a piecewise affine mixing function. Moran et al. (2021) imposed a restriction on how many observed variables are influenced by a single independent component, thus leading to a special kind of "sparsity" of the mixing function. In a slightly different context, Donoho and Grimes (2003) consider a constraint of orthogonality of the Jacobian (but for non-invertible \mathbf{f}) in the case of dimension reduction, which Horan et al. (2021) applied on nonlinear ICA combined with dimension reduction. Furthermore, Willetts and Paige (2021) automatically learn the auxiliary variables from observations by solving a secondary task, even without temporal structure; while Gresele et al. (2020b) used a different "view" of the data instead of auxiliary variable.

6 Nonlinear structural equation model

Just like the theory of linear ICA helped in solving the linear SEM problem in Sect. 4, it turns out that the theory of nonlinear ICA helps in solving the nonlinear SEM problem. In what follows we focus on the bivariate causal discovery problem. This corresponds to recovering the causal structure using observations from two variables, which we denote by x_1 and x_2 . While bivariate causal discovery is simplified special case of the more general causal discovery problem, it remains a challenging task.

6.1 Definition of problem

We start by defining a fully nonlinear SEM of arbitrary dimension as

$$x_j = f_j(\mathbf{PA}_j, e_j), \quad j = 1, \dots, n, \quad (23)$$

for arbitrary nonlinear functions f_j . The variables $\mathbf{PA}_j \subseteq \{x_1, \dots, x_n\} \setminus x_j$ are the parents of the variable x_j in the associated graph. The point is that the “parents” cause the variable x_j , and in the general case the problem of causal discovery can be framed as finding which variables are parents of which.

In order to accomplish identifiability, causal discovery algorithms generally adopt one of two techniques. The first approach is to impose constraints on the functions f_j that define the SEM (23). In the extreme case, we obtain linear models, possibly with non-Gaussian external influences as reviewed above. In contrast to the ICA case, however, some quite meaningful restricted nonlinear models have also been proposed for SEM and will be reviewed below. Nevertheless, arguably the most interesting case is where the nonlinearities are unrestricted while we impose some statistical assumptions on the external influences, which is our main topic here.

6.2 Identifiable weakly nonlinear causal models

In this subsection, we will review some of the most notable identifiable nonlinear causal models based on strong restrictions on the nonlinearity.

Additive noise model (ANM). Hoyer et al. (2009) introduced the additive noise model, in which the SEM has the form

$$x_j = f_j(\mathbf{PA}_j) + e_j,$$

and the noise variables \mathbf{e} are both mutually independent and independent of \mathbf{x} . Their theoretical identifiability result focuses on the case of two variables x_1 and x_2 . It stipulates that if x_1 causes x_2 , which we denote by $x_1 \rightarrow x_2$, then we cannot write $x_1 = g(x_2) + \tilde{e}$ for some function g and noise \tilde{e} that is independent of x_2 . Essentially, this SEM is asymmetrical with respect to x_1 and x_2 and can only describe the natural

cause–effect relationship. In other words, it is identifiable. Peters et al. (2014) generalized the identifiability result to the case of more than two variables.

Post-nonlinear model (PNL). Zhang and Hyvärinen (2009) introduced the post-nonlinear model, which generalizes ANM by adding a subsequent invertible mapping g_j :

$$x_j = g_j(f_j(\mathbf{PA}_j) + e_j).$$

The noise variables \mathbf{e} are still assumed to be mutually independent and independent of the causes. The authors show that the bivariate PNL model is identifiable in most cases and enumerate five special situations in which the model is not identifiable. This identifiability theory generalizes that of ANM, which is a special case when g_j is the identity mapping. Note that if we knew g_j , we could reduce the PNL model to an ANM by transforming the effect through the inverse of the mapping g_j , the transformed variable $g_j^{-1}(x_j)$ being a deterministic function of the original effect x_j .

Causal autoregressive flow (CAREFL). Khemakhem et al. (2021) note that SEMs are closely related to a class of model called (affine autoregressive) normalizing flows (Rezende and Mohamed 2015; Huang et al. 2018) in machine learning. That theory leads to the following definition of a SEM on the observations \mathbf{x} :

$$x_j = e^{\alpha_j(\mathbf{PA}_j)} z_j + \beta_j(\mathbf{PA}_j), \quad j = 1, 2 \tag{24}$$

where z_1, z_2 are statistically independent latent noise variables, and $\alpha_j(\mathbf{PA}_j)$ and $\beta_j(\mathbf{PA}_j)$ are scalar-valued functions, defined constant (with respect to \mathbf{x}) when there are no parents. This affine causal model generalizes the additive noise model (Hoyer et al. 2009) by adding a cause-dependent coefficient to the noise variable z_j in the SEM; thus, the disturbance $e^{\alpha_j(\mathbf{PA}_j)} z_j$ is not independent of the cause $\beta_j(\mathbf{PA}_j)$, in contrast to most models. (We use here the notation by Khemakhem et al. (2021) where z is used instead of e , the two terms on the RHS are switched, and Greek letters are used for the functions; but the real difference to ANM is in the modulation of the noise.) Identifiability of the model is shown by Khemakhem et al. (2021) under very general conditions. Recent work by Immer et al. (2022); Strobl and Lasko (2022) has modified and generalized the identifiability conditions.

6.3 Identifiability theory for general nonlinearities

Next we describe a method for estimating a nonlinear SEM for general nonlinearities; this leads to an identifiability result as well. We assume we observe bivariate data $\mathbf{x}(t) \in \mathbb{R}^2$ where t provides an index over all observations (t may be, e.g., a time index but this is not necessary).

Importantly, we further assume data is available over a set of distinct environmental conditions or segments, $C \in \mathcal{C}$. As such, each $\mathbf{x}(t)$ is allocated to one such environmental condition $C(t) \in \mathcal{C}$. In the case of causal discovery, such environmental conditions could be due to different interventions or due to some changes in the internal dynamics of the system being observed. To clearly align the proposed method with the terminology of nonlinear ICA reviewed in Sect. 5, we note that we

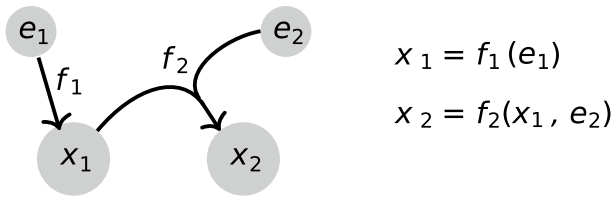


Fig. 4 Visualization of DAG, \mathcal{G} , associated with the SEM in Eqs. (25) and (26). The associated structural equations are provided on the right. Note that in contrast to Fig. 2, the disturbance variables are shown as well

may consider each environmental condition as a distinct segment in nonstationarity-based nonlinear ICA. As noted above, the segment should be interpreted in a very general sense, as in the theory of nonlinear ICA by auxiliary variables (Hyvärinen et al. 2019; Khemakhem et al. 2020a).

Next we outline the method for causal discovery over bivariate data by Monti et al. (2019), which they termed nonlinear SEM estimation using nonstationarity (NonSENS). Without loss of generality, we explain the basic logic assuming that $x_1 \rightarrow x_2$, such that the associated SEM is of the form:

$$x_1(t) = f_1(e_1(t)), \quad (25)$$

$$x_2(t) = f_2(x_1(t), e_2(t)), \quad (26)$$

where e_1, e_2 are latent disturbances whose distributions are assumed to vary across environmental conditions, thus creating the nonstationarity. The DAG associated with Eqs. (25) and (26) is shown in Fig. 4. The key idea here is that the latent disturbances in a bivariate SEM, \mathbf{e} , can be interpreted as corresponding to the independent sources in a nonlinear ICA model, \mathbf{s} , not unlike in the original LiNGAM theory reviewed above based on (Shimizu et al. 2006).

The proposed NonSENS algorithm consists of a two-step procedure. First, it seeks to recover latent disturbances via nonstationarity-based nonlinear ICA. We note that the nonstationarity introduced by the various environmental conditions, $C \in \mathcal{C}$, implies that the method by Hyvärinen and Morioka (2016) is well suited to recover the source variables e_1 and e_2 . Given estimated components, we may employ knowledge regarding the statistical independences between observed data and estimated sources in order to infer the causal structure, based on an interesting independence property of the components pointed out by Monti et al. (2019). Denote by $x \perp\!\!\!\perp y$ statistical independence of x and y , and by $\not\perp\!\!\!\perp$ lack of independence. We have:

Proposition 1 *Assume the true causal structure follows Eqs. (25) and (26), as depicted in Fig. 4. Then, it follows that $x_1 \perp\!\!\!\perp e_2$ while $x_1 \not\perp\!\!\!\perp e_1$ and $x_2 \not\perp\!\!\!\perp e_1$ as well as $x_2 \not\perp\!\!\!\perp e_2$.*

This proposition highlights the relationship between observations \mathbf{x} and true latent sources, \mathbf{e} : The cause is independent of one of the disturbances, while the effect is not independent of any disturbance. In practice, nonlinear ICA returns estimated latent sources whose ordering is random. As a result, we must test for independence between each of the variables, x_1 and x_2 , and each of the estimated latent sources, \hat{e}_1 and \hat{e}_2 . This results in a total of four distinct independence tests.

The proposed method is thus recapitulated as follows: After estimating (point-wise transformations of) latent variables by nonlinear ICA, we make the four statistical independence tests, and conclude a causal effect in the case where, for one of the directions, there is evidence to reject the null hypothesis of independence in three of the tests and only one of the tests fails to reject the null. In such a case, the observed variable regarding which one null hypothesis was not rejected is identified as the cause variable.

We finally note that the proposition and method just described based on Monti et al. (2019) constitute, in fact, a constructive identifiability proof: The direction of effect in arbitrary nonlinear SEM is identifiable, essentially assuming the data is divided into some conditions that fulfill the assumptions of nonlinear ICA methods based on nonstationarity.

7 Discussion

In this section, we discuss the actual utility of identifiability in applications, some questions for future research, as well as the question of algorithm development.

7.1 Utility of identifiability

7.1.1 Interpretability; finding causal direction

The utility of identifiability is obvious in the case where the parameters (or latent variables) of the model give directly some useful information about the phenomenon being analyzed. This is very clear in the case of causal analysis, where the whole point is to find out which variable causes which, and if the model is not identifiable, the analysis is typically not possible at all. Likewise, in the case of linear ICA, the components obtained often correspond to meaningful phenomena, as in the case of blind source separation (Fig. 1). This has been extensively used in neuroscience, for example, where the sources may correspond to sources of activity in the brain (Hyvärinen et al. 2010).

7.1.2 Feature extraction; semi-supervised and transfer learning

What is less understood is the utility of identifiability in nonlinear ICA. In deep learning, the components are often very difficult to interpret, since the nonlinearities implemented by the neural networks are difficult to understand or visualize. Therefore, the main utility of identifiability of nonlinear ICA might be found elsewhere. In fact, a representation learned by unsupervised deep learning is often used for some further purpose; for example, the features are used in a further classification task. So, the question is whether the identifiability of the components is useful for such a task.

To begin with, we note that even if the computation of the components is not possible to interpret, single components might still be useful for classification and decision making. For example, single components might work as biomarkers in a biomedical setting, and finding such biomarkers clearly requires identifiability. In a more general setting, Zhu et al. (2023) showed that nonlinear ICA is very useful for analyzing neuroscience data in a *semi-supervised setting*, which means that first the components are learned from a big generic data set, and then used on a smaller data set to solve a specific classification task using a simple linear classifier. In a similar vein, Khemakhem et al. (2020b) proposed some theory on how identifiable representations would be particularly useful for *transfer learning*, which is a related task where we have data from many different data sets (e.g., subjects in a biomedical setting) and then want to generalize to new subjects. While these developments point out the empirical utility of nonlinear ICA, they do not quite conclusively prove the utility of identifiability of nonlinear ICA; more developments on that topic would be warranted.

7.2 Questions for future research

7.2.1 Combining causal analysis with feature learning

In this paper, the methods are either finding some hidden factors or doing causal discovery. A very interesting topic would be to combine the two, so that the model learns hidden factors and causal connections *between* the hidden factors. This is what some would call “causal representation learning.” In the linear case, such methods can be found in Zhang and Hyvärinen (2009); Monti and Hyvärinen (2018). In the nonlinear case, this is a topic of great current interest in deep learning, see, e.g., Lachapelle et al. (2022); Morioka and Hyvärinen (2023) for some developments.

7.2.2 Dependent components; causal discovery with confounders

In fact, models with causal dependencies between the latent variables are a special case of models where the latent variables are not independent. A general framework for this case was proposed by Khemakhem et al. (2020b), who developed an extension of nonlinear ICA, called independently modulated component analysis (IMCA), where the components are allowed to dependent. The key idea is that the dependencies are assumed to be stationary, or independent of the auxiliary variable. Allowing such dependencies is likely to be useful in many contexts, and another highly promising topic for future research.

Another potential application of the IMCA framework in causal discovery is to allow for the presence of confounding. As already mentioned, a confounder is a hidden variable that affects both dependent and independent variables, resulting in spurious associations in the causal graph. We can establish a relationship between IMCA and confounded structural equation models (SEM) since the dependencies of the disturbances can be considered to be induced by the unobserved confounders. The identifiability of IMCA implies that the causal direction of such an SEM is likewise identifiable. However, because it is based on independence tests, the estimation technique of NonSENS based on testing described above cannot be used here. Possibly, a non-constraint-based method, such as likelihood ratio measures (Hyvärinen and Smith 2013; Monti et al. 2019), might still be pursued.

7.2.3 Identifiability of intermediate layers

If the nonlinear function in nonlinear ICA is modeled by a neural network, we are estimating much more than only a single nonlinearity. In fact, the intermediate layers in neural networks are frequently used as useful features for a later classification task. They may even be preferable in some applications over the representations learned by the final layer (Alain and Bengio 2018; Chen et al. 2020; Mikolov et al. 2013). An intriguing question arises: can the identifiability results of the representations learnt by nonlinear ICA be generalized to previous layers? Khemakhem et al. (2020b) showed that some such neural network architectures are, in fact, identifiable, using a form of induction to “propagate identifiability” forward through the network. Thus, a potential avenue of research is to prove that the intermediate layers preceding a final layer in a neural network can inherit the identifiability guarantees of the final layer. Extending such proofs to convolutional networks would be particularly interesting since they are frequently utilized in image learning and have a strong mathematical theory (Wiatowski and Bölcskei 2017).

7.3 Estimation algorithms

Finally, let us very briefly discuss what kind of estimation methods are available. Basically, we only need methods to estimate ICA, including nonlinear ICA, since the estimation of SEMs can be reduced to ICA. In machine learning, we typically

distinguish between two parts of an estimation method: 1) an estimation principle (such as maximum likelihood), typically leading to an objective function, and 2) a computational algorithm, typically optimizing the objective function. For linear ICA, we proposed FastICA (Hyvärinen 1999) which is widely used and was originally also used to solve the SEM estimation by Shimizu et al. (2006), although further methods for SEM were subsequently developed, e.g., by Shimizu et al. (2011); Hyvärinen and Smith (2013).

For nonlinear ICA, we have made several proposals; for an in-depth exposition, see Hyvärinen et al. (2023). First, we proposed “self-supervised” methods (Hyvärinen and Morioka 2016, 2017; Hyvärinen et al. 2019), which is a class of methods of great current interest in machine learning. Subsequently, we proposed maximum likelihood estimation, starting with models with additive noise (Khemakhem et al. 2020a; Hälvä et al. 2021) that enable variational approximation. The case of maximum likelihood estimation in the basic noise-free model used in this paper initially looks easy, but presents serious computational problems, which we have largely solved in Gresele et al. (2020a). In any case, developing better algorithms is definitely an interesting topic for future research as well.

Once we have an estimating algorithm, the question of its finite-sample performance can be considered. For linear ICA, the asymptotic variance (statistical efficiency) has been analyzed by several authors, including Cardoso and Laheld (1996); Pham and Garrat (1997); Hyvärinen (1997); Tichavsky et al. (2006), and these results may be more or less directly applicable to linear SEM estimation as well. Some analysis of the robustness against outliers has also been performed, in both linear and nonlinear cases (Hyvärinen 1997; Sasaki et al. 2020). For the nonlinear case, we are not aware of any analysis of statistical efficiency except in terms of simulations. These are clearly interesting points for future research.

8 Conclusion

We started this review from the linear, Gaussian factor analysis model which is known to be unidentifiable since the 1950’s if not earlier. The identifiability problem was solved by independent component analysis, a *non*-Gaussian factor analysis model which is identifiable in the precise sense that it can recover components that actually created the data, as in the case of blind source separation. Linear ICA further enables estimation of linear structural equation models, thus leading to identifiable causal discovery. ICA can be made nonlinear, but this is not at all straightforward: new assumptions are needed for identifiability. Here, we considered mainly

temporal dependencies and nonstationarity (the latter being taken in a very general sense). Restricting the nonlinear mixing is also likely to work but few successful models have been proposed so far. As in the linear case, we finally saw that nonlinear ICA enables identifiability and estimation of nonlinear SEM as well.

A. Liouville's theorem and nonlinear ICA

A possible line of research for making nonlinear ICA identifiable is imposing suitable conditions on the Jacobian of the transformation. The conditions are typically related to orthogonality of the Jacobian, also called local isometry (Gresele et al. 2021; Zimmermann et al. 2021; Buchholz et al. 2022). In this appendix, we consider the set of invertible mappings within a real space, i.e., such that the dimension is not changed, as is typical in ICA theory. We show that functions which are locally isometric and $\mathbb{R}^n \rightarrow \mathbb{R}^n$ are necessarily affine, and thus do not provide a meaningful basis for nonlinear ICA theory. We do this by proving a variant of the theorem on conformal mappings by Liouville. While the result can be considered well known, we provide a very simple proof for this variant, which is difficult to find in the literature.

A.1 Definitions

We start by some definitions:

Definition 1 Let \mathbf{f} be a differentiable mapping from an open subset U of \mathbb{R}^n to an open subset V of \mathbb{R}^n . The mapping called **conformal** if

$$\mathbf{Jf}(\mathbf{x})^T \mathbf{Jf}(\mathbf{x}) = c(\mathbf{x})\mathbf{I} \quad (27)$$

where \mathbf{Jf} is the Jacobian matrix (of partial derivatives) of \mathbf{f} , and $c(\mathbf{x})$ takes positive scalar values.

The geometrical meaning of the definition is that the mapping preserves angles, but the scaling can be changed by $c(\mathbf{x})$. An important special case is obtained when the Jacobian is orthogonal:

Definition 2 If in Def. 1, $c(\mathbf{x}) \equiv 1$, the mapping is called **locally isometric**.

In a geometrical interpretation, the mapping is required to preserve both angles and the scaling. This is related to the definition used by Donoho and Grimes (2003), and based on them, by Horan et al. (2021). However, those authors consider the case where the dimension is reduced by the mapping \mathbf{f} , and thus our results are rather different. In fact, local isometry is typically defined in the

context of manifolds which are of a lower dimension than the space itself, and therefore, our definition is a special case of the more conventional one. The orthogonality of the Jacobian has also been used, more heuristically, as a regularizer (Wei et al. 2021; Kumar and Poole 2020).

Merely for simplicity, we further introduce the following terminology:

Definition 3 A function \mathbf{f} is called **orthogonally affine** if it is of the form

$$\mathbf{f}(\mathbf{x}) = \mathbf{U}\mathbf{x} + \mathbf{b} \quad (28)$$

for some constant vector \mathbf{b} and an orthogonal matrix \mathbf{U} .

Such a function is also called “rigid motion” in some contexts.

A.2 Locally isometric functions are orthogonally affine

Our analysis is based on a well-known theorem on conformal mappings by Liouville from 1850. To keep the presentation short, we simply provide our variant of the theorem, and refer the reader to the literature for the original theorem (Flanders 1966; Nevanlinna 1960). Our variant of the theorem is as follows:

Theorem 2 Assume \mathbf{f} is locally isometric (Def. 2), and in \mathcal{C}^2 (i.e., with two continuous derivatives), in a real space of dimension $n \geq 2$. Then, \mathbf{f} is orthogonally affine (Def. 3).

Proof This proof is closely related to the first half of the proof of Liouville’s theorem by Flanders (1966); see Nevanlinna (1960) for another well-known proof.

We drop the argument \mathbf{x} for notational simplicity, and denote by $\mathbf{w} \cdot \mathbf{v}$ the dot-product. Denote by \mathbf{J}_i the i -th column of the Jacobian of \mathbf{f} , i.e., the vector of the partial derivatives with respect to x_i . Consider the i, j -th element of the matrix equation defining isometry by orthogonality of the Jacobian:

$$\mathbf{J}_i \cdot \mathbf{J}_j = \delta_{i=j} \quad (29)$$

Take the derivative with respect to x_k of both sides:

$$\mathbf{J}_i \cdot \mathbf{J}_j^k + \mathbf{J}_j \cdot \mathbf{J}_i^k = 0 \quad (30)$$

where \mathbf{J}_i^k is the vector of the partial derivatives with respect to x_k of the entries of \mathbf{J}_i . Thus, we get the following skew-symmetry condition:

$$\mathbf{J}_i \cdot \mathbf{J}_j^k = -\mathbf{J}_j \cdot \mathbf{J}_i^k. \quad (31)$$

On the other hand, by the fact that the order of derivation can be changed, we have the symmetricity condition:

$$\mathbf{J}_i \cdot \mathbf{J}_j^k = \mathbf{J}_i \cdot \mathbf{J}_k^j. \quad (32)$$

Obviously, a matrix which is both symmetric and skew-symmetric is necessarily zero. Here, we actually have a tensor but we show next that the same principle applies to such a tensor; this result is sometimes known as the Braid Lemma. Take any indices i, j, k , and apply (31) and (32) in alternation, three times. We get

$$\mathbf{J}_i \cdot \mathbf{J}_j^k = -\mathbf{J}_j \cdot \mathbf{J}_i^k = -\mathbf{J}_j \cdot \mathbf{J}_k^i = \mathbf{J}_k \cdot \mathbf{J}_j^i = \mathbf{J}_k \cdot \mathbf{J}_i^j = -\mathbf{J}_i \cdot \mathbf{J}_k^j = -\mathbf{J}_i \cdot \mathbf{J}_j^k. \quad (33)$$

The equality of the first and last term shows that they must be zero. This holds for all i , and the \mathbf{J}_i form an orthogonal basis, which implies that \mathbf{J}_j^k is zero. Thus, all second-order partial derivatives vanish at every point, and \mathbf{f} must be affine. By definition of isometry, it must further be orthogonally affine. \square

Our theorem extends the Liouville theory in the sense that our theorem applies to $n \geq 2$ while Liouville's theorem assumes $n \geq 3$. On the other hand, ours is a special case since Liouville assumes conformality while we assume local isometry. Given Liouville's theorem, it would be very easy to prove a corollary which gives the desired result for $n \geq 3$ but we prefer to prove a theorem which is not a strict corollary but an extension at the same time, without any help from such advanced theory. It was in fact possible to give a very simple proof in our case, while only the very complicated proofs seem to be available for Liouville's theorem.

The implication is that a nonlinear ICA model (or any deep latent variable model) where the (dimension-preserving) mixing is assumed locally isometric trivially reduces to a linear ICA model. (Note that we did not use any statistical properties of any components here.) The assumption of local isometry is far too strong from this viewpoint.

References

- Alain, G., Bengio, Y. (2018). Understanding intermediate layers using linear classifier probes. arXiv preprint [arXiv:1610.01644](https://arxiv.org/abs/1610.01644).
- Belouchrani, A., Meraim, K. A., Cardoso, J. F., Moulines, E. (1997). A blind source separation technique based on second order statistics. *IEEE Transactions on Signal Processing*, 45(2), 434–444.
- Bollen, K.A. (1989). *Structural Equations with Latent Variables*. Wiley.
- Brookes, M., Woolrich, M., Luckhoo, H., Price, D., Hale, J., Stephenson, M., Barnes, G., Smith, S., Morris, P. (2011). Investigating the electrophysiological basis of resting state networks using magnetoencephalography. *Proceedings of the National Academy of Sciences (USA)*, 108, 16783–16788.
- Buchholz, S., Besserve, M., Schölkopf, B. (2022). Function classes for identifiable nonlinear independent component analysis. arXiv preprint [arXiv:2208.06406](https://arxiv.org/abs/2208.06406).
- Cardoso, J. F. (2001). The three easy routes to independent component analysis: contrasts and geometry. *Proceedings of the International Workshop on Independent Component Analysis and Blind Signal Separation (ICA2001)*, San Diego.
- Cardoso, J. F., Laheld, B. H. (1996). Equivariant adaptive source separation. *IEEE Transactions on Signal Processing*, 44(12), 3017–3030.
- Chen, T., Kornblith, S., Norouzi, M., Hinton, G. (2020). A simple framework for contrastive learning of visual representations. arXiv preprint [arXiv:2002.05709](https://arxiv.org/abs/2002.05709).
- Comon, P. (1994). Independent component analysis—a new concept? *Signal Processing*, 36, 287–314.

- Donoho, D. L., Grimes, C. (2003). Hessian eigenmaps: Locally linear embedding techniques for high-dimensional data. *Proceedings of the National Academy of Sciences*, 100(10), 5591–5596.
- Donoho, D. L., Stodden, V. (2004). When does non-negative matrix factorization give a correct decomposition into parts? *Advances in Neural Information Processing 16 (Proceedings of NIPS2003)*. MIT Press.
- Eriksson, J., Koivunen, V. (2004). Identifiability, separability, and uniqueness of linear ICA models. *Signal Processing Letters, IEEE* 11(7), 601–604 (2004).
- Flanders, H. (1966). Liouville's theorem on conformal mapping. *Journal of Mathematics and Mechanics*, 15(1), 157–161.
- Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., Bengio, Y. (2014). Generative adversarial nets. *Advances in Neural Information Processing Systems*, 2672–2680.
- Gresele, L., Fissore, G., Javaloy, A., Schölkopf, B., Hyvärinen, A. (2020a). Relative gradient optimization of the jacobian term in unsupervised deep learning. *Advances in Neural Information Processing Systems (NeurIPS2020)*.
- Gresele, L., Rubenstein, P. K., Mehrjou, A., Locatello, F., Schölkopf, B. (2020b). The Incomplete Rosetta Stone Problem: Identifiability Results for Multi-View Nonlinear ICA. *Uncertainty in Artificial Intelligence*, 217–227. Proceedings of Machine Learning Research.
- Gresele, L., Von Kügelgen, J., Stümper, V., Schölkopf, B., Besserve, M. (2021). Independent mechanism analysis, a new concept? *Advances in neural information processing systems*, 34, 28233–28248.
- Hälvä, H., Hyvärinen, A. (2020). Hidden Markov nonlinear ICA: Unsupervised learning from nonstationary time series. *Proceedings of the 36th Conference on Uncertainty in Artificial Intelligence (UAI2020)*, Toronto.
- Hälvä, H., Corff, S. L., Lehericy, L., So, J., Zhu, Y., Gassiat, E., Hyvärinen, A. (2021). Disentangling identifiable features from noisy data with structured nonlinear ICA. *Advances in Neural Information Processing Systems (NeurIPS2021)*.
- Harman, H. H. (1967). *Modern Factor Analysis*. 2nd ed. University of Chicago Press.
- Harmeling, S., Ziehe, A., Kawanabe, M., Müller, K. R. (2003). Kernel-based nonlinear blind source separation. *Neural Computation*, 15(5), 1089–1124.
- Horan, D., Richardson, E., Weiss, Y. (2021). When is unsupervised disentanglement possible? *Advances in Neural Information Processing Systems*, 34, 5150–5161.
- Hoyer, P. O. (2004). Non-negative matrix factorization with sparseness constraints. *Journal of Machine Learning Research*, 5, 1457–1469.
- Hoyer, P. O., Janzing, D., Mooij, J., Peters, J., & Schölkopf, B. (2009). Nonlinear causal discovery with additive noise models. *Advances in Neural Information Processing Systems*, vol 21, pp 689–696. MIT Press.
- Hoyer, P. O., Shimizu, S., Kerminen, A. J., Palviainen, M. (2008). Estimation of causal effects using linear non-gaussian causal models with hidden variables. *International Journal of Approximate Reasoning*, 49, 362–378.
- Huang, C. W., Krueger, D., Lacoste, A., Courville, A. (2018). Neural autoregressive flows. *International Conference on Machine Learning*, 2078–2087. Proceedings of Machine Learning Research.
- Hyttinen, A., Barin-Pacela, V., Hyvärinen, A. (2022). Binary independent component analysis: A non-stationarity-based approach. *Proceedings of the 38th Conference on Uncertainty in Artificial Intelligence (UAI2022)*, 874–884, Eindhoven.
- Hyvärinen, A. (1997). One-unit contrast functions for independent component analysis: A statistical analysis. *Neural Networks for Signal Processing VII (Proceedings of the IEEE Workshop on Neural Networks for Signal Processing)*, 388–397, Amelia Island.
- Hyvärinen, A. (1999). Fast and robust fixed-point algorithms for independent component analysis. *IEEE Transactions on Neural Networks*, 10(3), 626–634.
- Hyvärinen, A., Morioka, H. (2016). Unsupervised feature extraction by time-contrastive learning and nonlinear ICA. *Advances in Neural Information Processing Systems (NIPS2016)*, Barcelona.
- Hyvärinen, A., Morioka, H. (2017). Nonlinear ICA of temporally dependent stationary sources. *Proceedings of the Artificial Intelligence and Statistics (AISTATS2017)*, Fort Lauderdale.
- Hyvärinen, A., Oja, E. (2000). Independent component analysis: Algorithms and applications. *Neural Networks*, 13(4–5), 411–430.
- Hyvärinen, A., Pajunen, P. (1999). Nonlinear independent component analysis: Existence and uniqueness results. *Neural Networks*, 12(3), 429–439.

- Hyvärinen, A., Smith S.M. (2013). Pairwise likelihood ratios for estimation of non-Gaussian structural equation models. *Journal of Machine Learning Research*, 14, 111–152.
- Hyvärinen, A., Karhunen, J., Oja, E. (2001). *Independent Component Analysis*. Wiley Interscience.
- Hyvärinen, A., Hurri, J., Hoyer, P. O. (2009). *Natural Image Statistics*. Springer-Verlag.
- Hyvärinen, A., Ramkumar, P., Parkkonen, L., Hari, R. (2010). Independent component analysis of short-time Fourier transforms for spontaneous EEG/MEG analysis. *NeuroImage*, 49(1), 257–271.
- Hyvärinen, A., Sasaki, H., Turner, R. (2019). Nonlinear ICA using auxiliary variables and generalized contrastive learning. *Proceedings of the Artificial Intelligence and Statistics (AISTATS2019)*, Okinawa.
- Hyvärinen, A., Khemakhem, I., Morioka, H. (2023). Nonlinear independent component analysis for principled disentanglement in unsupervised deep learning. arXiv preprint [arXiv:2303.16535](https://arxiv.org/abs/2303.16535).
- Immer, A., Schultheiss, C., Vogt, J. E., Schölkopf, B., Bühlmann, P., Marx, A. (2022). On the identifiability and estimation of causal location-scale noise models. arXiv preprint [arXiv:2210.09054](https://arxiv.org/abs/2210.09054).
- Jakobsen, M. E., Shah, R. D., Bühlmann, P., Peters, J. (2022). Structure learning for directed trees. *Journal of Machine Learning Research*, 23, 159.
- Jutten, C., Héroult, J. (1991). Blind separation of sources, part I: An adaptive algorithm based on neuro-mimetic architecture. *Signal Processing*, 24, 1–10.
- Khemakhem, I., Kingma, D. P., Monti, R. P., Hyvärinen, A. (2020a). Variational autoencoders and nonlinear ICA: A unifying framework. *Proceedings of the Artificial Intelligence and Statistics (AISTATS2020)*.
- Khemakhem, I., Monti, R. P., Kingma, D. P., Hyvärinen, A. (2020b). ICE-BeeM: Identifiable conditional energy-based deep models based on nonlinear ICA. *Advances in Neural Information Processing Systems (NeurIPS2020)*.
- Khemakhem, I., Monti, R. P., Leech, R., Hyvärinen, A. (2021). Causal autoregressive flows. *Proceedings of the Artificial Intelligence and Statistics (AISTATS2021)*.
- Kingma, D. P., Welling, M. (2014). Auto-encoding variational Bayes. *Proceedings of the International Conference on Learning Representations (ICLR2014)*, Banff.
- Kivva, B., Rajendran, G., Ravikumar, P., Aragam, B. (2022). Identifiability of deep generative models under mixture priors without auxiliary information. arXiv preprint [arXiv:2206.10044](https://arxiv.org/abs/2206.10044).
- Klindt, D., Schott, L., Sharma, Y., Ustyuzhaninov, I., Brendel, W., Bethge, M., Paiton, D. (2020). Towards nonlinear disentanglement in natural data with temporal sparse coding. arXiv preprint [arXiv:2007.10930](https://arxiv.org/abs/2007.10930).
- Kumar, A., Poole, B. (2020). On implicit regularization in β -vae. In *International Conference on Machine Learning*, 5480–5490. Proceedings of Machine Learning Research.
- Lacera, G., Spirtes, P., Ramsey, J., Hoyer, P. O. (2008). Discovering cyclic causal models by independent components analysis. *Proceedings of the 24th Conference on Uncertainty in Artificial Intelligence (UAI2008)*, Helsinki.
- Lachapelle, S., Rodriguez, P., Sharma, Y., Everett, K. E., Le Priol, R., Lacoste, A., Lacoste-Julien, S. (2022). Disentanglement via mechanism sparsity regularization: A new principle for nonlinear ICA. *Conference on Causal Learning and Reasoning*, 428–484. Proceedings of Machine Learning Research.
- Lee, D. D., Seung, H. S. (1999). Learning the parts of objects by non-negative matrix factorization. *Nature*, 401, 788–791.
- Locatello, F., Bauer, S., Lucic, M., Raetsch, G., Gelly, S., Schölkopf, B., Bachem, O. (2019). Challenging common assumptions in the unsupervised learning of disentangled representations. *International Conference on Machine Learning*, 4114–4124. Proceedings of Machine Learning Research.
- Matsuoka, K., Ohya, M., Kawamoto, M. (1995). A neural net for blind separation of nonstationary signals. *Neural Networks*, 8(3), 411–419.
- Mikolov, T., Chen, K., Corrado, G., Dean, J. (2013). Efficient Estimation of Word Representations in Vector Space. arXiv preprint [arXiv:1301.3781](https://arxiv.org/abs/1301.3781).
- Monti, R. P., Hyvärinen, A. (2018). A unified probabilistic model for learning latent factors and their connectivities from high-dimensional data. *Proceedings of the 34th Conference on Uncertainty in Artificial Intelligence (UAI2018)*, Monterey.
- Monti, R. P., Zhang, K., Hyvärinen, A. (2019). Causal discovery with general non-linear relationships using non-linear ICA. *Proceedings of the 35th Conference on Uncertainty in Artificial Intelligence (UAI2019)*, Tel Aviv.
- Moran, G. E., Sridhar, D., Wang, Y., Blei, D. M. (2021). Identifiable variational autoencoders via sparse decoding. arXiv preprint [arXiv:2110.10804](https://arxiv.org/abs/2110.10804).

- Morioka, H., Hälvä, H., Hyvärinen, A. (2021). Independent innovation analysis for nonlinear vector autoregressive process. *Proceedings of the Artificial Intelligence and Statistics (AISTATS2021)*.
- Morioka, H., Hyvärinen, A. (2023). Connectivity-contrastive learning: Combining causal discovery and representation learning for multimodal data. *Proceedings of the Artificial Intelligence and Statistics (AISTATS2023)*, Valencia, Spain.
- Nevanlinna, R. (1960). On differentiable mappings. *Analytic functions*, 3–9.
- Olshausen, B. A., Field, D.J. (1997). Sparse coding with an overcomplete basis set: A strategy employed by V1? *Vision Research*, 37, 3311–3325.
- Paatero, P., Tapper, U. (1994). Positive matrix factorization: A non-negative factor model with optimal utilization of error estimates of data values. *Environmetrics*, 5, 111–126.
- Pearl, J. (2009). *Causality*, Cambridge University Press.
- Peters, J., Bühlmann, P. (2014). Identifiability of gaussian structural equation models with equal error variances. *Biometrika*, 101(1), 219–228.
- Peters, J., Janzing, D., Schölkopf, B. (2017). *Elements of causal inference: foundations and learning algorithms*. MIT press.
- Peters, J., Mooij, J. M., Janzing, D., Schölkopf, B. (2014). Causal discovery with continuous additive noise models. *Journal of Machine Learning Research*, 15, 2009–2053.
- Pham, D. T., Cardoso, J. F. (2001). Blind separation of instantaneous mixtures of nonstationary sources. *IEEE Transactions Signal Processing*, 49(9), 1837–1848.
- Pham, D. T., Garrat, P. (1997). Blind separation of mixture of independent sources through a quasi-maximum likelihood approach. *IEEE Transactions on Signal Processing*, 45(7), 1712–1725.
- Plumbley, M.D. (2003). Algorithms for non-negative independent component analysis. *IEEE Transactions on Neural Networks*, 14(3), 534–543.
- Rezende, D. J., Mohamed, S. (2015). Variational inference with normalizing flows. arXiv preprint [arXiv:1505.05770](https://arxiv.org/abs/1505.05770).
- Sasaki, H., Takenouchi, T., Monti, R. P., Hyvärinen, A. (2020). Robust contrastive learning and nonlinear ICA in the presence of outliers. *Proceedings of the 36th Conference on Uncertainty in Artificial Intelligence (UAI2020)*, Toronto.
- Schell, A., Oberhauser, H. (2023). Nonlinear independent component analysis for discrete-time and continuous-time signals. *Annals of Statistics*, In press.
- Shimizu, S. (2014). LiNGAM: Non-Gaussian methods for estimating causal structures. *Behaviormetrika*, 41(1), 65–98.
- Shimizu, S., Hoyer, P. O., Hyvärinen, A., Kerminen, A. (2006). A linear non-Gaussian acyclic model for causal discovery. *Journal of Machine Learning Research*, 7, 2003–2030.
- Shimizu, S., Inazumi, T., Sogawa, Y., Hyvärinen, A., Kawahara, Y., Washio, T., Hoyer, P. O., Bollen, K. (2011). DirectLiNGAM: A direct method for learning a linear non-gaussian structural equation model. *Journal of Machine Learning Research*, 12, 1225–1248.
- Spirtes, P., Glymour, C., Scheines, R., Heckerman, D., Meek, C., Richardson, T. (2000). *Causation, Prediction and Search*, MIT Press.
- Spirtes, P., Zhang, K. (2016). Causal discovery and inference: concepts and recent methodological advances. *Applied Informatics*.
- Sprekeler, H., Zito, T., Wiskott, L. (2014). An extension of slow feature analysis for nonlinear blind source separation. *Journal of Machine Learning Research*, 15(1), 921–947.
- Strobl, E. V., Lasko, T. A. (2022). Identifying patient-specific root causes with the heteroscedastic noise model. arXiv preprint [arXiv:2205.13085](https://arxiv.org/abs/2205.13085).
- Tashiro, T., Shimizu, S., Hyvärinen, A., Washio, T. (2014). ParceLiNGAM: a causal ordering method robust against latent confounders. *Neural Computation*, 26, 57–83.
- Tichavsky, P., Koldovsky, Z., Oja, E. (2006). Performance analysis of the fastica algorithm and crame/spl acute/r-rao bounds for linear independent component analysis. *IEEE Transactions on Signal Processing*, 54(4), 1189–1203.
- Tong, L., Liu, R. W., Soon, V. C., Huang, Y. F. (1991). Indeterminacy and identifiability of blind identification. *IEEE Transactions on Circuits and Systems*, 38, 499–509.
- Wei, Y., Shi, Y., Liu, X., Ji, Z., Gao, Y., Wu, Z., Zuo, W. (2021). Orthogonal Jacobian regularization for unsupervised disentanglement in image generation. *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 6721–6730.

- Wiatowski, T., Bölcskei, H. (2017). A Mathematical Theory of Deep Convolutional Neural Networks for Feature Extraction. arXiv preprint [arXiv:1512.06293](https://arxiv.org/abs/1512.06293).
- Willets, M., Paige, B. (2021). I Don't Need u: Identifiable Non-Linear ICA Without Side Information. arXiv preprint [arXiv:2106.05238](https://arxiv.org/abs/2106.05238).
- Zhang, K., Chan, L. (2008). Minimal nonlinear distortion principle for nonlinear independent component analysis. *Journal of Machine Learning Research*, 9, 2455–2487.
- Zhang, K., Hyvärinen, A. (2009). On the identifiability of the post-nonlinear causal model. *Proceedings of the 25th Conference on Uncertainty in Artificial Intelligence (UAI2009)*, 647–655, Montréal.
- Zhang, K., Hyvärinen, A. (2010). Source separation and higher-order causal analysis of MEG and EEG. In *Proceedings of the 26th Conference on Uncertainty in Artificial Intelligence (UAI2010)*, Catalina Island.
- Zhu, Y., Parviainen, T., Heinilä, E., Parkkonen, L., Hyvärinen, A. (2023). Unsupervised representation learning of spontaneous MEG data with nonlinear ICA. *NeuroImage*. In press.
- Zimmermann, R. S., Sharma, Y., Schneider, S., Bethge, M., Brendel, W. (2021). Contrastive learning inverts the data generating process. *International Conference on Machine Learning*, 12979–12990. Proceedings of Machine Learning Research.

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Springer Nature or its licensor (e.g. a society or other partner) holds exclusive rights to this article under a publishing agreement with the author(s) or other rightsholder(s); author self-archiving of the accepted manuscript version of this article is solely governed by the terms of such publishing agreement and applicable law.