

Structural Causal 3D Reconstruction

Weiyang Liu^{1,2*}, Zhen Liu^{3*}, Liam Paull³, Adrian Weller^{2,4}, and Bernhard Schölkopf¹

¹Max Planck Institute for Intelligent Systems, Tübingen

²University of Cambridge ³Mila, Université de Montréal ⁴Alan Turing Institute

Abstract. This paper considers the problem of unsupervised 3D object reconstruction from in-the-wild single-view images. Due to ambiguity and intrinsic ill-posedness, this problem is inherently difficult to solve and therefore requires strong regularization to achieve disentanglement of different latent factors. Unlike existing works that introduce explicit regularizations into objective functions, we look into a different space for implicit regularization – the structure of latent space. Specifically, we restrict the structure of latent space to capture a topological causal ordering of latent factors (*i.e.*, representing causal dependency as a directed acyclic graph). We first show that different causal orderings matter for 3D reconstruction, and then explore several approaches to find a task-dependent causal factor ordering. Our experiments demonstrate that the latent space structure indeed serves as an implicit regularization and introduces an inductive bias beneficial for reconstruction.

1 Introduction

Understanding the 3D structures of objects from their 2D views has been a longstanding and fundamental problem in computer vision. Due to the lack of high-quality 3D data, unsupervised single-view 3D reconstruction is typically favorable; however, it is an ill-posed problem by nature, and it typically requires a number of carefully-designed priors and regularizations to achieve good disentanglement of latent factors [31, 61, 7, 33, 6, 75, 15]. Distinct from these existing works that focus on introducing explicit regularizations, we aim to explore how the structure of latent space can implicitly regularize 3D reconstruction, and to answer the following question: *Can a suitable structure of latent space encode helpful implicit regularization and yield better inductive bias?*

Current single-view 3D reconstruction methods [75, 33, 39] typically decompose 3D objects into several semantic latent factors such as 3D shape, texture, lighting and viewpoint. These latent factors are independently extracted from single 2D images and then fed into a differentiable renderer to reconstruct the original 2D images, as illustrated in Fig. 1(a). Conditioned on the input image, these latent factors are typically assumed to be independent from each other. Such an assumption for disentanglement can be too strong and sometimes unrealistic, because it suggests that the estimated viewpoint will not affect the estimation of lighting in the image, which contradicts the formation of realistic images. This observation motivates us to explore how the dependency structure of latent factors implicitly regularizes the encoder and improves disentanglement.

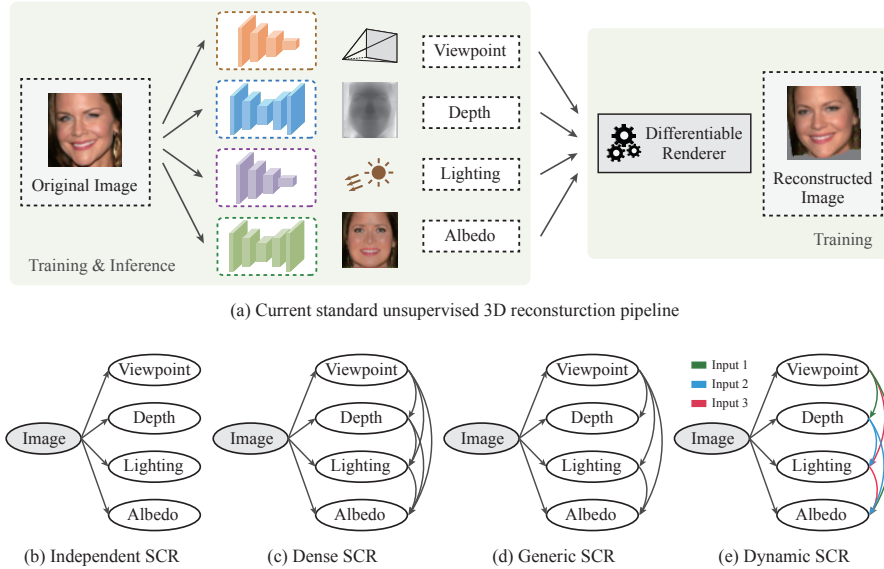


Fig. 1: (a) Overview of the standard 3D reconstruction pipeline. (b) Graphical model of independent SCR, which is adopted in the standard pipeline and assumes full conditional independence. (c) Graphical model of a dense SCR example. This makes no assumption on the distribution. (d) Graphical model of a generic SCR example. This assumes partial conditional independence. (e) Graphical model of a dynamic SCR example. This yields strong flexibility. All directed edges are learned in practice.

Taking inspiration from structural causal models [53], we propose the **Structural Causal Reconstruction (SCR)** framework which introduces structural priors to the latent space. We consider the causal ordering of latent factors and study how different causal orderings can introduce different inductive biases.

Depending on the type of causal orderings and the corresponding flexibility, we derive three SCR variants: dense SCR which learns a chain factorization without any embedded conditional independence, generic SCR which learns a directed acyclic graph (DAG) over the latent factors, and dynamic SCR which learns a dynamic DAG that is dependent on the input image. We note that the standard 3D reconstruction pipeline can be viewed as independent SCR as shown in Fig. 1(b) (*i.e.*, viewpoint, depth, lighting and albedo are conditionally independent from each other given the input image), while dense SCR does not assume any conditional independence. Generic SCR learns a DAG over the latent factors and serves as an interpolation between independence SCR and dense SCR by incorporating partial conditional independence. Both dense SCR and generic SCR are learned with a static ordering which is fixed once trained. To accommodate the over-simplified rendering model and the complex nature of image formation, we propose dynamic SCR that can capture more complex dependency by learning input-dependent DAGs. This can be useful when modeling in-the-wild images that are drawn from a complex multi-modal distribution [47]. Specifically, we apply Bayesian optimization to dense SCR to search for the best

dense causal ordering of the latent factors. For generic SCR, we first propose to directly learn a DAG with an additional regularization. Besides that, we further propose a two-phase algorithm: first running dense SCR to obtain a dense ordering and then learning the edges via masking. For dynamic SCR, we propose a self-attention approach to learn input-dependent DAGs.

From a distribution perspective, independent SCR (Fig. 1(b)) is the least expressive graphical model in the sense that it imposes strong conditional independence constraints and therefore limits potential distributions that can factorize over it. On the contrary, any conditional distribution $P(\mathbf{V}, \mathbf{D}, \mathbf{L}, \mathbf{A} | \mathbf{I})$ (where $\mathbf{V}, \mathbf{D}, \mathbf{L}, \mathbf{A}, \mathbf{I}$ denote viewpoint, depth, lighting, albedo and image, respectively) can factorize over dense SCR, making it the most expressive variant for representing distributions. Generic SCR unifies both independent SCR and dense SCR by incorporating a flexible amount of conditional independence constraints. Dynamic SCR is able to capture even more complex conditional distribution that is dynamically changing for different input images.

Intuition for why learning a latent dependency structure helps 3D reconstruction comes from the underlying entanglement among estimated viewpoint, depth, lighting and albedo. For example, conditioned on a given 2D image, a complete disentanglement between viewpoint and lighting indicates that changing the estimated viewpoint of an object will not change its estimated lighting. This makes little sense, since changing the viewpoint will inevitably affect the estimation of lighting. In contrast to existing pipelines that extract viewpoint and depth independently from the image (Fig. 2(a)), the information of viewpoint may give constraints on lighting (*e.g.*, modeled as a directed edge from \mathbf{V} to \mathbf{L} in Fig. 2(b)). Therefore, instead of ignoring the natural coupling among latent factors and assuming conditional independence, we argue that learning a suitable dependency structure for latent factors is crucial for intrinsic disentanglement. In general, modeling the latent dependency and causality among viewpoint, depth, lighting and albedo renders an implicit regularization for disentanglement, leading to strong generalizability. Beside the intuition from the anti-causal direction, Section 3.3 gives another interpretation for SCR from the causal direction. Our contributions are:

- We explicitly model the causal structure among the latent factors.
- To learn a causal ordering, we propose three SCR variants including dense SCR, generic SCR and dynamic SCR. Each one yields a different level of distribution expressiveness and modeling flexibility.
- We constrain the latent space structure to be a topological causal ordering (which can represent arbitrary DAGs), reducing the difficulty of learning.
- Our method is in parallel to most current 3D reconstruction pipelines and can be used simultaneously with different pipelines such as [75, 62, 15, 39].
- Our empirical results show that different causal orderings of latent factors lead to significantly different 3D reconstruction performance.

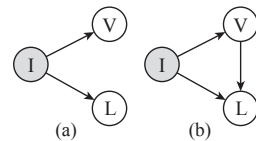


Fig. 2: (a) Viewpoint and lighting are extracted independently from the input image. (b) The extracted viewpoint gives constraints on lighting. These arrows denote encoding latent variables from the image (*i.e.*, anti-causal direction).

2 Related Work

Multi-view 3D reconstruction. This method usually requires multi-view images of the same target object. Classical techniques such as Structure from Motion [49] and Simultaneous Localization and Mapping [18] rely on hand-crafted geometric features and matching across different views. Owing to the availability of large 3D object datasets, modern approaches [9, 32, 77] can perform multi-view 3D reconstruction with neural networks that map 2D images to 3D volumes.

Shape from X. There are many alternative monocular cues that can be used for reconstructing shapes from images, such as shading [28, 82], silhouettes [36], texture [74] and symmetry [46, 16]. These methods are generally not applicable to in-the-wild images due to their strong assumptions. Shape-from-symmetry [46, 16, 63, 60] assumes the symmetry of the target object, making use of the original image and its horizontally flipped version as a stereo pair for 3D reconstruction. [60] demonstrates the possibility to detect symmetries and correspondences using descriptors. Shape-from-shading assumes a specific shading model (*e.g.*, Phong shading [54] and spherical harmonic lighting [22]), and solves an inverse rendering problem to decompose different intrinsic factors from 2D images.

Single-view 3D reconstruction. This line of research [9, 21, 23, 64, 73, 78, 86, 13, 25, 39, 75, 29, 12] aims to reconstruct a 3D shape from a single-view image. [68, 50, 71] use images and their corresponding ground truth 3D meshes as supervisory signals. This, however, requires either annotation efforts [76] or synthetic construction [5]. To avoid 3D supervision, [34, 41, 33, 7] consider an analysis-by-synthesis approach with differentiable rendering, but they still require either multi-view images or known camera poses. To further reduce supervision, [31] learns category-specific 3D template shapes from an annotated image collection, but annotated 2D keypoints are still necessary in order to infer camera pose correctly. [26] also studies a similar category-specific 3D reconstruction from a single image. [39] estimates 3D mesh, texture and camera pose of both rigid and non-rigid objects from a single-view image using silhouette as supervision. Videos [1, 85, 48, 66, 72] are also leveraged as a form of supervision for single-view 3D reconstruction. For human bodies and faces, [30, 20, 67, 19, 80, 8, 15, 14, 72, 4] reconstruct 3D shapes from single-view images with a predefined shape model such as SMPL [45], FLAME [38] or BFM [52]. Among many works in single-view 3D reconstruction, we are particularly interested in a simple and generic unsupervised framework from [75] that utilizes the symmetric object prior. This framework adopts the Shape-from-shading pipeline to extract intrinsic factors of images, including 3D shape, texture, viewpoint and illumination parameters (as shown in Fig. 1(a)). The encoders are trained to minimize the reconstruction error between the input image and the rendered image. It shows impressive results in reconstructing human faces, cat faces and synthetic cars.

For the sake of simplicity, we build the SCR pipeline based on the framework of [75] and focus on studying how the causal structure of latent factors affects the 3D reconstruction performance. We emphasize that our method is a parallel contribution to [75] and is generally applicable to any 3D reconstruction framework without the need of significant modifications.

3 Causal Ordering of Latent Factors Matters

The very first question we need to address is “*Does the causal ordering of latent factors matter for unsupervised 3D reconstruction?*”. Without an affirmative answer, it will be pointless to study how to learn a good causal ordering.

3.1 A Motivating Example from Function Approximation

We start with a motivating example to show the advantages of modeling the dependency between latent factors. We take a look at the example in Fig. 3 where the lighting factor \mathbf{L} can be represented using either $f_L(\mathbf{I})$ in Fig. 3(a) or $f_L(\mathbf{I}) + h_L(\mathbf{V})$ in Fig. 3(b). There are a few perspectives to compare these two representations and see their difference (also see Appendix C):

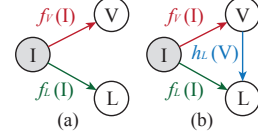


Fig. 3: Two structures encode the lighting factor.

- We first assume the underlying data generating function for lighting is given by $\mathbf{L} := f_L^*(\mathbf{I}) + h_L^*(\mathbf{V})$ where f_L^* and h_L^* are two polynomial functions of order p . Because $\mathbf{V} := f_V^*(\mathbf{I})$ where f_V^* is also a polynomial function of order p , we can then write the lighting function as $\mathbf{L} := f_L^*(\mathbf{I}) + h_L^* \circ f_V^*(\mathbf{I})$ which is a polynomial function order $2p$. The lighting function can be learned with either $\mathbf{L} = f_L(\mathbf{I})$ in Fig. 3(a) or $\mathbf{L} = f_L(\mathbf{I}) + h_L \circ f_V(\mathbf{I})$ in Fig. 3(b). The previous requires the encoder $f_L(\mathbf{I})$ to learn a polynomial of order $2p$, while the latter requires learning that of only order p .
- From the perspective of function approximation, it is obvious that $f_L(\mathbf{I}) + h_L \circ f_V(\mathbf{I})$ is always more expressive than $f_L(\mathbf{I})$ given that f_L, h_L, f_V are of the same representation capacity. Therefore, the structure shown in Fig. 3(b) is able to capture more complex and nonlinear lighting function.
- Making the lighting \mathbf{L} partially dependent on the viewpoint \mathbf{V} gives the lighting function an inherent structural prior, which may implicitly regularizes the function class and constrain its inductive bias.

3.2 Expressiveness of Representing Conditional Distributions

The flexibility of SCR can also be interpreted from a distribution perspective. Most existing 3D reconstruction pipelines can be viewed as independent SCR whose conditional distribution $P(\mathbf{V}, \mathbf{D}, \mathbf{L}, \mathbf{A} | \mathbf{I})$ can be factorized into

$$P(\mathbf{V}, \mathbf{D}, \mathbf{L}, \mathbf{A} | \mathbf{I}) = P(\mathbf{V} | \mathbf{I}) \cdot P(\mathbf{D} | \mathbf{I}) \cdot P(\mathbf{L} | \mathbf{I}) \cdot P(\mathbf{A} | \mathbf{I}) \quad (1)$$

which renders the conditional independence among $\mathbf{V}, \mathbf{D}, \mathbf{L}, \mathbf{A}$. This is in fact a strong assumption that largely constrains the potential family of distributions that can factorize over this model, making this model less expressive in representing conditional distributions. In contrast, dense SCR does not assume any conditional independence because it yields the following factorization (this is just one of the potential orderings and we randomly choose one for demonstration):

$$P(\mathbf{V}, \mathbf{D}, \mathbf{L}, \mathbf{A} | \mathbf{I}) = P(\mathbf{V} | \mathbf{I}) \cdot P(\mathbf{D} | \mathbf{I}, \mathbf{V}) \cdot P(\mathbf{L} | \mathbf{I}, \mathbf{V}, \mathbf{D}) \cdot P(\mathbf{A} | \mathbf{I}, \mathbf{V}, \mathbf{D}, \mathbf{L}) \quad (2)$$

which imposes no constraints to the factorized conditional distribution and is more expressive. Therefore, any dense ordering has this nice property of assuming

no conditional independence among latent factors. However, there exists a trade-off between expressiveness and learnability. A more expressive model usually requires more data to train and is relatively sample-inefficient. Generic SCR is proposed in search of a sweet spot between expressiveness and learnability by incorporating partial conditional independence. Taking Fig. 1(c) as an example, we can observe that this model assumes $P(\mathbf{D} \perp \mathbf{L}|\mathbf{I})$ and $P(\mathbf{D} \perp \mathbf{A}|\mathbf{I})$. Going beyond generic SCR, dynamic SCR aims to tackle with the scenario where the conditional distribution $P(\mathbf{V}, \mathbf{D}, \mathbf{L}, \mathbf{A}|\mathbf{I})$ is dynamically changing rather than being static for all the images. This can greatly enhance the modeling flexibility.

3.3 Modeling Causality in Rendering-based Decoding

The previous subsection shows that there is no difference for different dense orderings in representing $P(\mathbf{V}, \mathbf{D}, \mathbf{L}, \mathbf{A}|\mathbf{I})$. This conclusion is drawn from the perspective of modeling correlation. However, one of the most significant properties of topological ordering is its ability to model acyclic causality. In terms of causal relationships, different orderings (including both dense and generic ones) make a difference. The standard 3D reconstruction pipeline is naturally an autoencoder architecture, where the encoder and decoder can be interpreted as anti-causal and causal mappings, respectively [56, 55, 70, 3, 35, 37]. Here, the causal part is a generative mapping, and the anti-causal part is in the opposite direction, inferring causes from effects. However, how to determine which part of the pipeline should be viewed as anti-causal or causal remains unclear. Here we discuss three possible partitions of anti-causal and causal mappings, as shown in Fig. 4. The partition denoted by green dashed line uses an identity mapping as the anti-causal direction and the rest of the pipeline performs causal reconstruction. This partition does not explicitly model the causes and may not be useful. For the partition labeled by the blue dashed line, all the encoders are viewed as anti-causal, so the latent factor ordering is also part of anti-causal learning and does not necessarily benefit from the underlying causal ordering (*i.e.*, causal DAG [65], cf. [37]). The partition denoted by the red dashed line views part of the encoder as anti-causal learning and the rest of the encoder along with the renderer as causal learning. This partition is particularly interesting because it puts the latent factor ordering to the causal direction and effectively connects latent factor ordering to the underlying causal ordering. Our SCR framework (in Section 4.1) is designed based on such insight. When the underlying causal ordering is available, us-

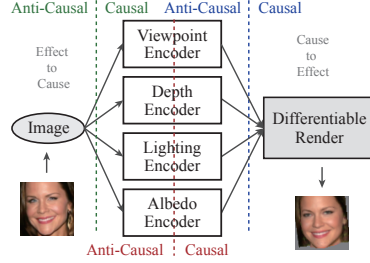


Fig. 4: Three possible partitions of anti-causal and causal mappings.

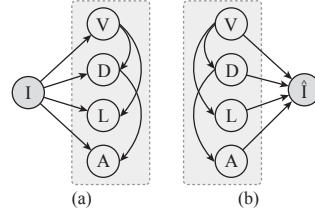


Fig. 5: Latent structure modeling from (a) anti-causal direction and (b) causal direction. Gray regions denote where the causal ordering is learned.

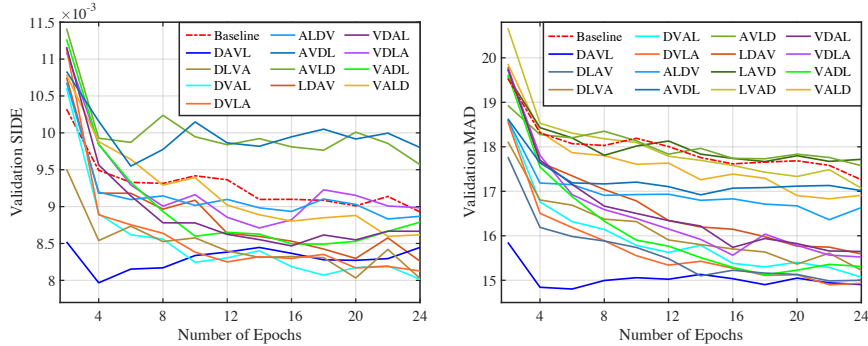


Fig. 6: The scale-invariant depth error (left) and mean angle deviation (right) on the BFM dataset [52] for different dense causal orderings. For visualization clarity, we plot the SIDE of the best three orderings, the worst three orderings and random six orderings. For MAD, we plot the same selection of orderings along with the best three and worst three orderings. We denote depth, albedo, lighting and viewpoint as D, A, L and V, respectively. For the full results, please refer to Appendix B.

ing it as the default ordering could be beneficial. Although the causal ordering could improve strong generalization [35], learning the causal ordering without additional knowledge (*e.g.*, interventions or manipulations such as randomized experiment) is difficult and out of our scope. *We hypothesize that the underlying causal ordering leads to fast, generalizable and disentangled 3D reconstruction, and learning causal ordering based on these criteria may help us identify crucial causal relations.* As an encouraging signal, one of the best-performing dense ordering (DAVL) well matches the conventional rendering procedures in OpenGL, which is likely to be similar to the underlying causal ordering.

In the previous examples of Fig. 2 and Fig. 3, we justify the necessity of the topological ordering from the factor estimation (*i.e.*, anti-causal) perspective. As discussed above, we can alternatively incorporate the causal ordering to the causal mapping and model the causality among latent factors in the decoding (*i.e.*, generative) process, which well matches the design of structural causal models. This is also conceptually similar to [79, 58] except that SCR augments the decoder with a physics-based renderer. Fig. 5 shows two interpretations of latent factor ordering from the causal and anti-causal directions. While the causal mapping encourages SCR to approximate the underlying causal ordering, the anti-causal mapping does not necessarily do so. The final learned causal ordering may be the result of a trade-off between causal and anti-causal mapping.

3.4 Empirical Evidence on 3D Reconstruction

Most importantly, we demonstrate the empirical performance of different dense causal orderings for unsupervised 3D reconstruction. The details of our pipeline and the experimental settings are given in Section 4.1 and Appendix A, respectively. Here we focus on comparing different dense orderings. As can be observed from Fig. 6, different settings for dense SCR yield significantly different empirical behaviors, validating our claim that topological causal ordering of latent factors matters in unsupervised 3D reconstruction. Moreover, we discover that most

of the dense orderings perform consistently for both SIDE and MAD metrics. For example, depth-albedo-viewpoint-lighting, depth-viewpoint-albedo-lighting and depth-viewpoint-lighting-albedo perform consistently better than the other dense orderings and the baseline (*i.e.*, independent SCR). This again matches our intuition in Section 3.3 that different dense ordering indicates different causality and leads to different disentanglement/reconstruction performance despite being equivalent in representing the conditional distribution $P(\mathbf{V}, \mathbf{D}, \mathbf{L}, \mathbf{A} | \mathbf{I})$.

Interestingly, the well-performing dense orderings also seem to match our knowledge about the underlying causal ordering. For example, we also tend to put viewpoint in front of lighting, because the viewpoint will cause the change of lighting effects on the object. Almost all the well-performing dense orderings have this pattern, suggesting that the well-performing orderings tend to match the intrinsic causality that is typically hard to obtain in practice.

4 Learning Causal Ordering for 3D Reconstruction

We introduce a generic framework to learn causal ordering. Our proposed pipeline and algorithms to learn different variants of SCR are by no means optimal ones and it remains an open problem to learn a good causal ordering. We instead aim to show that a suitable causal ordering is beneficial to 3D reconstruction.

4.1 General SCR Framework

Our unsupervised 3D reconstruction pipeline is inspired by [75] but with some novel modifications to better accommodate the learning of causal ordering. Our goal is to study how causal ordering affects the disentanglement and generalizability in 3D reconstruction rather than achieving state-of-the-art performance.

Decoding from a common embedding space. A differentiable renderer typically takes in latent factors of different dimensions, making it less convenient to incorporate causal factor ordering. In order to easily combine multiple latent factors, we propose a learnable decoding method that includes additional neural networks ($f_V^2, f_D^2, f_L^2, f_A^2$ shown in Fig. 7) to the differentiable renderer. These neural networks transform the latent factors from a common d -dimensional embedding space ($\mathbf{u}_V, \mathbf{u}_D, \mathbf{u}_L, \mathbf{u}_A$) to their individual dimensions ($\mathbf{V}, \mathbf{D}, \mathbf{L}, \mathbf{A}$) such that the differentiable renderer can directly use them as inputs.

Implementing SCR in a common embedding space. Since all the latent factors can be represented in a common embedding space of the same dimension, we now introduce how to implement SCR in this pipeline. We start by listing a few key desiderata: (1) all variants of SCR should have (roughly) the same number of trainable parameters as independent SCR (baseline) such that the comparison is meaningful; (2) learning SCR should be efficient, differentiable and end-to-end; (3) different structures among latent factors can be explored in a unified framework by imposing different constraints on the adjacency matrix.

We first interpret conditional probability in terms of neural networks. For example, $P(\mathbf{V} | \mathbf{I}, \mathbf{D})$ can be implemented as a single neural network $\mathbf{V} = f_V(\mathbf{I}, \mathbf{D})$ that takes both image \mathbf{I} and depth \mathbf{D} as input. Instead of parameterizing the encoder f_V with one neural network, we separate f_V into two neural networks

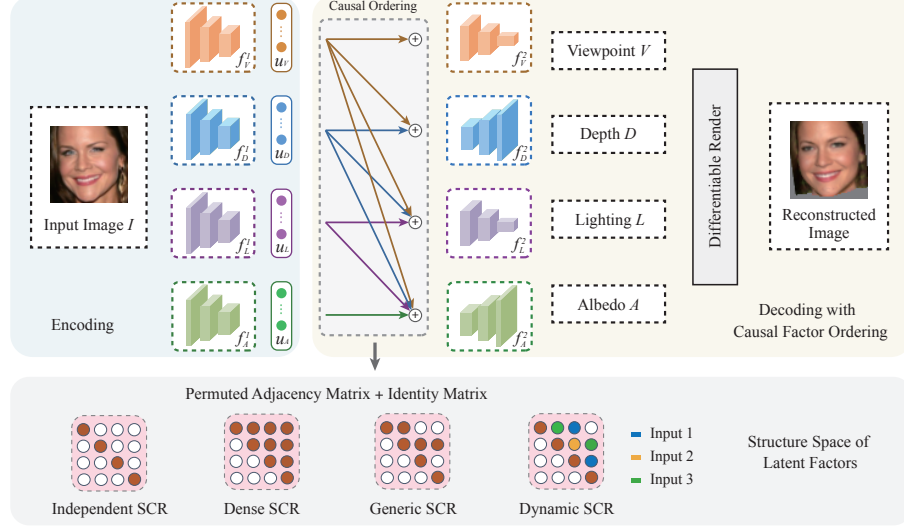


Fig. 7: Our unsupervised 3D reconstruction pipeline to explore causal ordering. The causal edges in the figure are for illustration. Actual edges are learned in practice.

f_V^1, f_V^2 – the first one f_V^1 aims to map different factors into a common embedding space of the same dimension, and the second one f_V^2 transforms the embedding to the final factor that can be used directly for the differentiable renderer. Taking $P(V|I, D)$ as an example, we model it using $V = f_V^2(f_V^1(I), u_D)$. We define the SCR adjacency matrix that characterizes the dependency structure among latent factors as $\mathbf{M} = [\mathbf{M}_V, \mathbf{M}_D, \mathbf{M}_L, \mathbf{M}_A] \in \mathbb{R}^{4 \times 4}$ where $\mathbf{M}_V = [M_{VV}, M_{VD}, M_{VL}, M_{VA}]^\top \in \mathbb{R}^{4 \times 1}$ and M_{VD} denotes the weight of the directed edge from V to D (the weight can be constrained to be either binary or continuous). Because causal ordering is equivalent to DAG, \mathbf{M} can be permuted into a strictly upper triangular matrix. Generally, latent factors are modeled by

$$\begin{aligned} V &= f_V^2(f_V^1(I), \mathbf{M}_V^\top \mathbf{u}) & D &= f_D^2(f_D^1(I), \mathbf{M}_D^\top \mathbf{u}) \\ L &= f_L^2(f_L^1(I), \mathbf{M}_L^\top \mathbf{u}) & A &= f_A^2(f_A^1(I), \mathbf{M}_A^\top \mathbf{u}) \end{aligned} \quad (3)$$

where $\mathbf{u} = [u_V; u_D; u_L; u_A] \in \mathbb{R}^{4 \times d}$. The input to $f_V^2, f_D^2, f_L^2, f_A^2$ can either be added element-wisely or concatenated, and we use element-wise addition in order not to introduce additional parameters. \mathbf{M} exactly implements causal ordering as an equivalent form of causal DAG. More generally, \mathbf{M} characterizes the latent space structure and can also be constrained to be some other family of structures.

Interpreting SCR as a part of causal mapping. After modeling the latent factors with two separate neural networks, we can view $f_V^1, f_D^1, f_L^1, f_A^1$ as the encoding process (*i.e.*, the light blue region in Fig. 7). Different from [75], we view the causal ordering, $f_V^2, f_D^2, f_L^2, f_A^2$ and differentiable renderer as the decoding process (*i.e.*, the light yellow region in Fig. 7). This can be understood as an augmented trainable physics-based renderer which performs rendering with additional neural networks and a causal ordering. More importantly, incorporating

causal ordering to the decoding process makes it a part of causal mapping, which may produce more interpretable ordering due to its intrinsic connection to the underlying causality. Therefore, our novel pipeline design makes it possible to benefit from (or even estimate) the underlying causal ordering.

Loss functions. To avoid introducing additional priors to SCR and better study the effect of causal ordering, we stick to the same loss functions as [75]. The loss function is defined as $\mathcal{L} = \mathcal{L}_{\text{rec}}(\hat{\mathbf{I}}, \mathbf{I}) + \lambda_f \mathcal{L}_{\text{rec}}(\hat{\mathbf{I}}', \mathbf{I}) + \lambda_p \mathcal{L}_p(\hat{\mathbf{I}}, \mathbf{I})$ where \mathcal{L}_{rec} is the reconstruction loss and \mathcal{L}_p is the perceptual loss. λ_f, λ_p are hyperparameters $\hat{\mathbf{I}}$ is the reconstructed image with original depth and albedo. $\hat{\mathbf{I}}'$ is the reconstructed image with flipped depth and albedo. Similar to [75], we also use the confidence map to compensate asymmetry. Appendix A provides the detailed formulation.

Learning causal ordering. We formulate the SCR learning as a bi-level optimization where the inner optimization is to train the 3D reconstruction networks with \mathcal{L} and the outer optimization learns a suitable adjacency matrix \mathbf{M} :

$$\min_{\mathbf{M} \in \mathcal{M}_{\text{DAG}}} \mathcal{L}_{\text{val}}(\mathbf{W}^*(\mathbf{M}), \mathbf{M}) \quad \text{s.t. } \mathbf{W}^*(\mathbf{M}) = \arg \min_{\mathbf{W}} \mathcal{L}_{\text{train}}(\mathbf{W}, \mathbf{M}) \quad (4)$$

where \mathbf{W} denotes all the trainable parameters of neural networks in the 3D reconstruction pipeline, including $f_V^1, f_V^2, f_D^1, f_D^2, f_L^1, f_L^2, f_A^1, f_A^2$. $\mathcal{L}_{\text{train}}$ is the loss \mathcal{L} computed on the training set, and \mathcal{L}_{val} is the loss \mathcal{L} computed on the validation set. Optionally, \mathcal{L}_{val} may also include other supervised losses (*e.g.*, ground truth depth) if available. This is in general a difficult problem, and in order to solve it effectively, we propose different algorithms based on the properties of the feasible set \mathcal{M}_{DAG} . After \mathbf{M} is learned, we will fix \mathbf{M} and retrain the network.

4.2 Learning Dense SCR via Bayesian Optimization

The adjacency matrix \mathbf{M} for dense SCR is an all-one strictly upper triangular matrix after proper permutation. Therefore, we are essentially learning the ordering permutation which is a discrete and non-differentiable structure. We resort to Bayesian optimization (BO) [17] that is designed for gradient-free and “expensive to evaluate” optimization. Specifically, BO first places a Gaussian process prior on $\mathcal{L}_{\text{val}}(\mathbf{W}^*(\mathbf{M}), \mathbf{M})$ in Eq. (4) and collect all the evaluated points on \mathbf{M} . Then BO updates posterior probability distribution on \mathcal{L}_{val} using all available data and evaluates \mathcal{L}_{val} on the maximizer point of the acquisition function which is computed with the current posterior distribution. Note that, evaluation on \mathcal{L}_{val} requires computing $\mathbf{W}^*(\mathbf{M})$. Finally, BO outputs the latest evaluated \mathbf{M} . We use the position permutation kernel $K(\pi_1, \pi_2 | \lambda) = \exp(-\lambda \cdot \sum_i |\pi_1^{-1}(i) - \pi_2^{-1}(i)|)$ where π is a permutation mapping that maps the original index to the permuted index. We use the expected improvement as the acquisition function. We note a special advantage of BO over gradient-dependent methods: the validation metric can be obtained from user study, which is often more reliable and flexible.

4.3 Learning Generic SCR via Optimization Unrolling

To solve the bi-level optimization in Eq. (4), we can unroll the inner optimization with a few gradient updates and replace $\mathbf{W}^*(\mathbf{M})$ with $\mathbf{W} - \eta \nabla_{\mathbf{W}} \mathcal{L}_{\text{train}}(\mathbf{W}, \mathbf{M})$. Then the optimization becomes $\min_{\mathbf{M} \in \mathcal{M}_{\text{DAG}}} \mathcal{L}_{\text{val}}(\mathbf{W} - \eta \nabla_{\mathbf{W}} \mathcal{L}_{\text{train}}(\mathbf{W}, \mathbf{M}), \mathbf{M})$.

Here we unroll 1-step gradient update as an example, but we can also unroll multiple steps for better performance in practice. In order to constrain the adjacency matrix \mathbf{M} to be a DAG, we can turn the feasible set $\mathbf{M} \in \mathcal{M}_{\text{DAG}}$ into a constraint [84]: $\mathcal{H}(\mathbf{M}) = \text{tr}((\mathbf{I}_n + \frac{c}{n}\mathbf{M} \circ \mathbf{M})^n) - n = 0$ where \mathbf{I}_n is an identity matrix of size n , c is some arbitrary positive number, n is the number of latent factors (here $n = 4$) and \circ denotes the element-wise multiplication. Using Lagrangian multiplier method, we end up with the following optimization:

$$\min_{\mathbf{M}} \mathcal{L}_{\text{val}}(\mathbf{W} - \eta \nabla_{\mathbf{W}} \mathcal{L}_{\text{train}}(\mathbf{W}, \mathbf{M}), \mathbf{M}) + \lambda_{\text{DAG}} \mathcal{H}(\mathbf{M}) \quad (5)$$

where λ_{DAG} is a hyperparameter. Alternatively, we may use the augmented Lagrangian method for stronger regularization [84, 81]. Although Eq. (5) is easy to optimize, it is still difficult to guarantee the learned \mathbf{M} to be a strict DAG and the search space may also be too large. To address this, we further propose a different approach to learn \mathbf{M} . The basic idea is to learn generic SCR based on the solution from dense SCR. We simply need to relearn/remove some edges for the given dense ordering. The final optimization is given by

$$\min_{\mathbf{M}} \mathcal{L}_{\text{val}}(\mathbf{W} - \eta \nabla_{\mathbf{W}} \mathcal{L}_{\text{train}}(\mathbf{W}, \mathbf{M} \circ \mathbf{M}_{\text{dense}}^*), \mathbf{M} \circ \mathbf{M}_{\text{dense}}^*) \quad (6)$$

where $\mathbf{M}_{\text{dense}}^*$ is obtained from BO for dense SCR. It is a binary matrix that can be permuted to be strictly upper triangular. If we also constrain \mathbf{M} to be binary, we will use a preset threshold to binarize the obtained \mathbf{M} before retraining.

4.4 Learning Dynamic SCR via Masked Self-Attention

In order to make the adjacency matrix \mathbf{M} be adaptively dependent on the input, we need to turn \mathbf{M} into the output of a function that takes the image \mathbf{I} as input, *i.e.*, $\mathbf{M} = \Phi(\mathbf{I})$. One sensible choice is to parameterize $\Phi(\cdot)$ with an additional neural network, but it will inevitably introduce significantly more parameters and increase the capacity of the framework, making it unfair to compare with the other variants. Therefore, we take a different route by utilizing self-attention to design $\Phi(\cdot)$. Specifically, we use $\mathbf{M} = \Phi(\mathbf{I}) = q(\mathbf{u}) \circ \mathbf{M}_{\text{dense}}^*$ where \mathbf{u} is the matrix containing all the factor embeddings ($\mathbf{u} = [\mathbf{u}_V; \mathbf{u}_D; \mathbf{u}_L; \mathbf{u}_A]$), $q(\mathbf{u})$ can be either the Sigmoid activation $\sigma(\mathbf{u}\mathbf{u}^\top)$ or cosine cross-similarity matrix among $\mathbf{u}_V, \mathbf{u}_D, \mathbf{u}_L, \mathbf{u}_A$ (*i.e.*, $q(\mathbf{u})_{i,j} = \frac{\langle \mathbf{u}_i, \mathbf{u}_j \rangle}{\|\mathbf{u}_i\| \|\mathbf{u}_j\|}$, $i, j \in \{V, D, L, A\}$), and $\mathbf{M}_{\text{dense}}^*$ is the solution obtained from BO for dense SCR. $\mathbf{M}_{\text{dense}}^*$ essentially serves as a mask for the self-attention such that the resulting causal ordering is guaranteed to be a DAG. Since there is no fixed \mathbf{M} , the entire pipeline is trained with the final objective function: $\min_{\mathbf{W}} \mathcal{L}_{\text{train}}(\mathbf{W}, q(\mathbf{u}) \circ \mathbf{M}_{\text{dense}}^*)$ in an end-to-end fashion. We note that the function $q(\mathbf{u})$ has no additional parameters and meanwhile makes the causal ordering (*i.e.*, $\Phi(\mathbf{I})$) dynamically dependent on the input image \mathbf{I} .

4.5 Insights and Discussion

Connection to neural architecture search. We discover an intriguing connection between SCR and neural architecture search (NAS) [87, 11, 40]. SCR can be viewed as a special case of NAS that operates on a semantically interpretable space (*i.e.*, the dependency structure among latent factors), while standard NAS does not necessarily produce an interpretable architecture. SCR performs like

a top-down NAS where a specific neural structure is derived from semantic dependency/causality and largely constrains the search space for neural networks without suffering from countless poor local minima like NAS does.

Semantic decoupling in common embeddings. In order to make SCR interpretable, we require the latent embeddings $\mathbf{u}_V, \mathbf{u}_D, \mathbf{u}_L, \mathbf{u}_A$ to be semantically decoupled. For example, \mathbf{u}_V should contain sufficient information to decode \mathbf{V} . The semantic decoupling in the common embedding space can indeed be preserved. First, the DAG constraint can naturally encourage semantic decoupling. We take an arbitrary dense ordering (*e.g.*, DAVL) as an example. \mathbf{u}_D is the only input for f_D^2 , so it contains sufficient information for \mathbf{D} . \mathbf{u}_D, \mathbf{A} are the inputs for f_A^2 , so the information of \mathbf{A} will be largely encoded in \mathbf{u}_A (\mathbf{u}_D already encodes the information of \mathbf{D}). The same reasoning applies to \mathbf{V} and \mathbf{L} . Note that, a generic DAG will have less decoupling than dense ordering due to less number of directed edges. Second, we enforce the encoders $f_V^2, f_D^2, f_L^2, f_A^2$ to be relatively simple functions (*e.g.*, shallow neural networks), such that they are unable to encode too much additional information and mostly serve as dimensionality transformation. They could also be constrained to be invertible. Both mechanisms ensure the semantic decoupling in the common embedding space.

5 Experiments and Results

Datasets. We evaluate our method on two human face datasets (CelebA [44] and BFM [52]), one cat face dataset that combines [83] and [51] (cropped by [75]) and one car dataset [75] rendered from ShapeNet [5] with random viewpoints and illumination. These images are split 8:1:1 into training, validation and testing.

Metrics. For fairness, we use the same metrics as [75]. The first one is Scale Invariant Depth Error (SIDE) [10] which computes the standard deviation of the difference between the estimated depth map at the input view and the ground truth depth map at the log scale. We note that this metric may not reflect the true reconstruction quality. As long as this metric is reasonably low, it may no longer be a stronger indicator for reconstruction quality, which is also verified by [27]. To make a comprehensive evaluation, we also use another metric: the mean angle deviation (MAD) [75] between normals computed from ground truth depth from the predicted depth. It measures how well the surface is reconstructed.

Implementation. For the network architecture, we follow [75] and only make essential changes to its setup such that the comparison is meaningful. For the detailed implementation and experimental settings, refer to Appendix A.

5.1 Quantitative Results

Geometry reconstruction. We train and test all the methods on the BFM dataset to evaluate the depth reconstruction quality. The results are given in Table 1. We compare different variants of SCR with our own baseline (*i.e.*, independent SCR), two state-of-the-art methods [75, 27], supervised learning upper bound, constant null depth and average ground truth depth. We note that there is a performance difference between our re-run version and the original version of [75]. This is because all our experiments are run under CUDA-10

while the original version of [75] is trained on CUDA-9. We also re-train our models on CUDA-9 and observe a similar performance boost (see Appendix D). We suspect this is because of the rendering precision on different CUDA versions. However, this will not affect the advantages of our method and our experiment settings are the same for all the other compared methods.

More importantly, we build our SCR on a baseline that performs similarly to [75] (independent SCR vs. our version of [75]). SCR improves our baseline for more than 0.0065 on SIDE and 2.5 degree on MAD. Specifically, our dense SCR learns an ordering of depth-viewpoint-albedo-lighting. Generic SCR (Eq. 6) and both dynamic SCR variants are built upon this ordering. We notice that if we use a random dense ordering, then the 3D reconstruction results are even worse than our baseline, which shows that dense SCR can indeed learn crucial structures. Such a significant performance gain shows that a suitable SCR can implicitly regularize the neural networks and thus benefit the 3D reconstruction.

Effect of dense ordering.

We also perform ablation study to see how generic SCR and dynamic SCR perform if they are fed with different dense orderings. Table 2 compares three different dense orderings: depth-light-viewpoint-albedo (DLVA), depth-albedo-viewpoint-lighting (DAVL) and depth-viewpoint-albedo-lighting (DVAL). We show that G-SCR and D-SCR can consistently improve the 3D reconstruction results even if different dense orderings are given as the mask.

Convergence. Fig. 8 plots the convergence curves of both SIDE and MAD in Fig. 8. We observe that dense SCR, generic SCR and dynamic SCR converge much faster than the baseline. Dense SCR achieves impressive performance at the very beginning of the training. When converged, dynamic SCR and generic SCR performs better than dense SCR due to its modeling flexibility.

Method	SIDE ($\times 10^{-2}$) \downarrow	MAD (deg.) \downarrow
Supervised	0.410 ± 0.103	10.78 ± 1.01
Constant Null Depth	2.723 ± 0.371	43.34 ± 2.25
Average GT Depth	1.990 ± 0.556	23.26 ± 2.85
Wu et al. [75] (reported)	0.793 ± 0.140	16.51 ± 1.56
Ho et al. [27] (reported)	0.834 ± 0.169	15.49 ± 1.50
Wu et al. [75] (our run)	0.901 ± 0.190	17.53 ± 1.84
Independent SCR	0.895 ± 0.183	17.36 ± 1.78
Dense SCR (random)	1.000 ± 0.275	17.66 ± 2.09
Dense SCR (BO)	0.830 ± 0.205	14.88 ± 1.94
Generic SCR (Eq. 5)	0.859 ± 0.215	15.17 ± 1.92
Generic SCR (Eq. 6)	0.820 ± 0.190	14.79 ± 1.96
Dynamic SCR (Sigmoid)	0.827 ± 0.220	14.86 ± 2.02
Dynamic SCR (Cosine)	0.815 ± 0.232	14.80 ± 1.95

Table 1: Depth reconstruction results on BFM.

Method	DLVA	DAVL	DVAL
Dense SCR (fixed)	15.02 ± 2.00	15.14 ± 1.91	14.88 ± 1.94
G-SCR (Eq. 6)	14.96 ± 1.90	14.85 ± 2.13	14.79 ± 1.96
Dy-SCR (Sigmoid)	15.01 ± 1.99	15.03 ± 2.12	14.86 ± 2.02
Dy-SCR (Cosine)	14.99 ± 1.93	15.05 ± 2.15	14.80 ± 1.95

Table 2: MAD (degree) results on BFM.

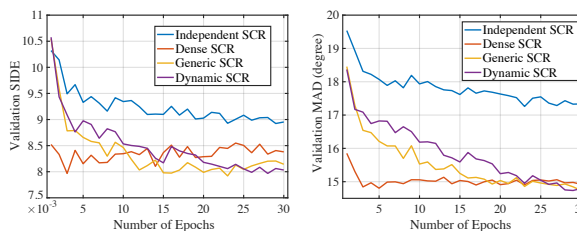


Fig. 8: Convergence curves of validation SIDE & MAD.

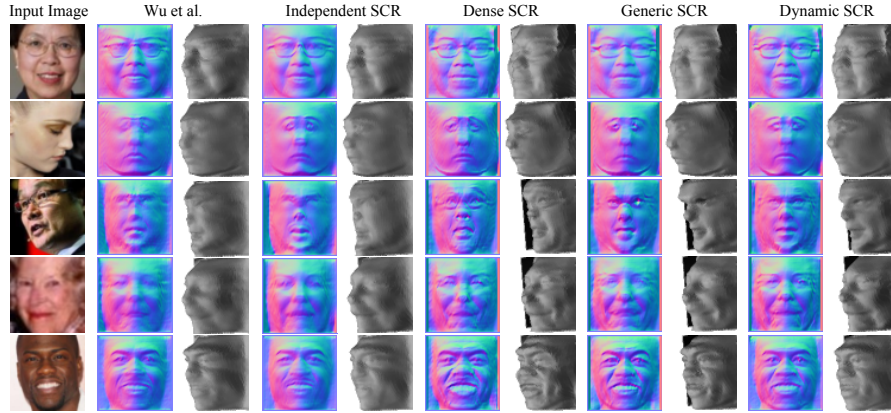


Fig. 9: One textureless view and canonical normal map on CelebA. All the methods (including Wu et al. [75]) are trained on CelebA under the same experimental settings.

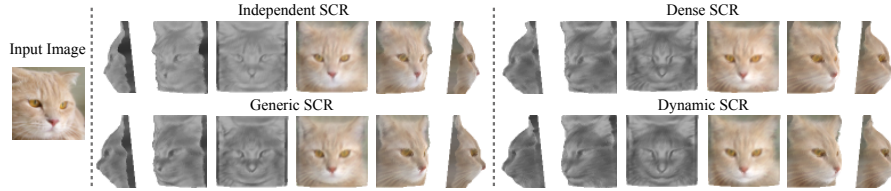


Fig. 10: Textured and textureless shapes from multiple views on cat faces.

5.2 Qualitative Results

CelebA. We show the reconstruction results for a few challenging in-the-wild face images (*e.g.*, extreme poses and expressions) in Fig. 9. We train dense SCR with BO and the other SCR variants are trained based on the best learned dense ordering. Our SCR variants including dense SCR, generic SCR and dynamic SCR are able to reconstruct fine-grained geometric details and recover more realistic shapes than both [75] and our independent baseline. This well verifies the importance of implicit regularization from latent space structure.

Cat faces. We also train all the SCR variants on cat faces. Results in Fig. 10 show that dynamic SCR yields the best 3D reconstruction quality, while dense SCR and generic SCR can also recover reasonably good geometric details.

Cars. We train all the methods on the synthetic car dataset under the same settings, and then evaluate these methods on car images with abundant geometric details. Fig. 11 shows that our SCR variants recovers very fine-grained geometric details and produce highly realistic normal maps which are significantly better than both [75] and independent SCR.

Conclusion. We demonstrate the potential of causality in regularizing 3D reconstruction. For acknowledgements, see <https://wyliu.com/papers/ECCV.txt>.

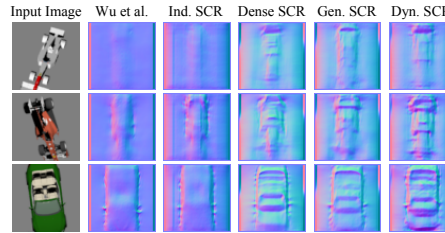


Fig. 11: Canonical normal maps on cars.

References

1. Agrawal, P., Carreira, J., Malik, J.: Learning to see by moving. In: ICCV (2015) 4
2. Albiero, V., Chen, X., Yin, X., Pang, G., Hassner, T.: img2pose: Face alignment and detection via 6dof, face pose estimation. In: CVPR (2021)
3. Besserve, M., Sun, R., Schölkopf, B.: Intrinsic disentanglement: an invariance view for deep generative models. In: ICML 2018 Workshop on Theoretical Foundations and Applications of Deep Generative Models (2018) 6
4. Blanz, V., Vetter, T.: A morphable model for the synthesis of 3d faces. In: Proceedings of the 26th annual conference on Computer graphics and interactive techniques (1999) 4
5. Chang, A.X., Funkhouser, T., Guibas, L., Hanrahan, P., Huang, Q., Li, Z., Savarese, S., Savva, M., Song, S., Su, H., et al.: Shapenet: An information-rich 3d model repository. arXiv preprint arXiv:1512.03012 (2015) 4, 12
6. Chen, C.H., Tyagi, A., Agrawal, A., Drover, D., Stojanov, S., Rehg, J.M.: Unsupervised 3d pose estimation with geometric self-supervision. In: CVPR (2019) 1
7. Chen, W., Ling, H., Gao, J., Smith, E., Lehtinen, J., Jacobson, A., Fidler, S.: Learning to predict 3d objects with an interpolation-based differentiable renderer. NeurIPS (2019) 1, 4
8. Choutas, V., Pavlakos, G., Bolkart, T., Tzionas, D., Black, M.J.: Monocular expressive body regression through body-driven attention. In: ECCV (2020) 4
9. Choy, C.B., Xu, D., Gwak, J., Chen, K., Savarese, S.: 3d-r2n2: A unified approach for single and multi-view 3d object reconstruction. In: ECCV (2016) 4
10. Eigen, D., Puhrsch, C., Fergus, R.: Depth map prediction from a single image using a multi-scale deep network. In: NeurIPS (2014) 12
11. Elsken, T., Metzen, J.H., Hutter, F.: Neural architecture search: A survey. The Journal of Machine Learning Research 20(1), 1997–2017 (2019) 11
12. Fahim, G., Amin, K., Zarif, S.: Single-view 3d reconstruction: a survey of deep learning methods. Computers & Graphics 94, 164–190 (2021) 4
13. Fan, H., Su, H., Guibas, L.J.: A point set generation network for 3d object reconstruction from a single image. In: CVPR (2017) 4
14. Feng, Y., Choutas, V., Bolkart, T., Tzionas, D., Black, M.J.: Collaborative regression of expressive bodies using moderation. In: 3DV (2021) 4
15. Feng, Y., Feng, H., Black, M.J., Bolkart, T.: Learning an animatable detailed 3d face model from in-the-wild images. ACM Transactions on Graphics (TOG) (2021) 1, 3, 4
16. François, A.R., Medioni, G.G., Waupotitsch, R.: Mirror symmetry = 2-view stereo geometry. Image and Vision Computing 21(2), 137–143 (2003) 4
17. Frazier, P.I.: A tutorial on bayesian optimization. arXiv preprint arXiv:1807.02811 (2018) 10
18. Fuentes-Pacheco, J., Ruiz-Ascencio, J., Rendón-Mancha, J.M.: Visual simultaneous localization and mapping: a survey. Artificial intelligence review 43(1), 55–81 (2015) 4
19. Gecer, B., Ploumpis, S., Kotsia, I., Zafeiriou, S.: Ganfit: Generative adversarial network fitting for high fidelity 3d face reconstruction. In: CVPR (2019) 4
20. Gerig, T., Morel-Forster, A., Blumer, C., Egger, B., Luthi, M., Schönborn, S., Vetter, T.: Morphable face models-an open framework. In: 2018 13th IEEE International Conference on Automatic Face & Gesture Recognition (FG 2018) (2018) 4

21. Girdhar, R., Fouhey, D.F., Rodriguez, M., Gupta, A.: Learning a predictable and generative vector representation for objects. In: ECCV (2016) 4
22. Green, R.: Spherical harmonic lighting: The gritty details. In: Archives of the game developers conference. vol. 56, p. 4 (2003) 4
23. Gwak, J., Choy, C.B., Chandraker, M., Garg, A., Savarese, S.: Weakly supervised 3d reconstruction with adversarial constraint. In: 3DV (2017) 4
24. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: CVPR (2016)
25. Henderson, P., Ferrari, V.: Learning to generate and reconstruct 3d meshes with only 2d supervision. arXiv preprint arXiv:1807.09259 (2018) 4
26. Henderson, P., Ferrari, V.: Learning single-image 3d reconstruction by generative modelling of shape, pose and shading. IJCV **128**(4), 835–854 (2020) 4
27. Ho, L.N., Tran, A.T., Phung, Q., Hoai, M.: Toward realistic single-view 3d object reconstruction with unsupervised learning from multiple images. In: ICCV (2021) 12, 13
28. Horn, B.K., Brooks, M.J.: Shape from shading. MIT press (1989) 4
29. Hu, T., Wang, L., Xu, X., Liu, S., Jia, J.: Self-supervised 3d mesh reconstruction from single images. In: CVPR (2021) 4
30. Kanazawa, A., Black, M.J., Jacobs, D.W., Malik, J.: End-to-end recovery of human shape and pose. In: CVPR (2018) 4
31. Kanazawa, A., Tulsiani, S., Efros, A.A., Malik, J.: Learning category-specific mesh reconstruction from image collections. In: ECCV (2018) 1, 4
32. Kar, A., Häne, C., Malik, J.: Learning a multi-view stereo machine. In: NIPS (2017) 4
33. Kato, H., Harada, T.: Learning view priors for single-view 3d reconstruction. In: CVPR (2019) 1, 4
34. Kato, H., Ushiku, Y., Harada, T.: Neural 3d mesh renderer. In: CVPR (2018) 4
35. Kilbertus, N., Parascandolo, G., Schölkopf, B.: Generalization in anti-causal learning. arXiv preprint arXiv:1812.00524 (2018) 6, 7
36. Koenderink, J.J.: What does the occluding contour tell us about solid shape? Perception **13**(3), 321–330 (1984) 4
37. Leeb, F., Lanzillotta, G., Annadani, Y., Besserve, M., Bauer, S., Schölkopf, B.: Structure by architecture: Disentangled representations without regularization. arXiv preprint arXiv:2006.07796 (2020) 6
38. Li, T., Bolkart, T., Black, M.J., Li, H., Romero, J.: Learning a model of facial shape and expression from 4d scans. ACM Transactions on Graphics **36**(6), 1–17 (2017) 4
39. Li, X., Liu, S., Kim, K., De Mello, S., Jampani, V., Yang, M.H., Kautz, J.: Self-supervised single-view 3d reconstruction via semantic consistency. In: ECCV (2020) 1, 3, 4
40. Liu, H., Simonyan, K., Yang, Y.: Darts: Differentiable architecture search. In: ICLR (2019) 11
41. Liu, S., Li, T., Chen, W., Li, H.: Soft rasterizer: A differentiable renderer for image-based 3d reasoning. In: ICCV (2019) 4
42. Liu, W., Wen, Y., Raj, B., Singh, R., Weller, A.: Sphereface revived: Unifying hyperspherical face recognition. TPAMI (2022)
43. Liu, W., Wen, Y., Yu, Z., Li, M., Raj, B., Song, L.: Sphereface: Deep hypersphere embedding for face recognition. In: CVPR (2017)
44. Liu, Z., Luo, P., Wang, X., Tang, X.: Deep learning face attributes in the wild. In: ICCV (2015) 12

45. Loper, M., Mahmood, N., Romero, J., Pons-Moll, G., Black, M.J.: Smpl: A skinned multi-person linear model. *ACM Transactions on Graphics* **34**(6), 1–16 (2015) 4
46. Mukherjee, D.P., Zisserman, A.P., Brady, M., Smith, F.: Shape from symmetry: Detecting and exploiting symmetry in affine images. *Philosophical Transactions of the Royal Society of London. Series A: Physical and Engineering Sciences* **351**(1695), 77–106 (1995) 4
47. Murphy, K.P.: *Dynamic bayesian networks: representation, inference and learning*. University of California, Berkeley (2002) 2
48. Novotny, D., Larlus, D., Vedaldi, A.: Learning 3d object categories by looking around them. In: *ICCV* (2017) 4
49. Ozyesil, O., Voroninski, V., Basri, R., Singer, A.: A survey of structure from motion. *arXiv preprint arXiv:1701.08493* (2017) 4
50. Pan, J., Han, X., Chen, W., Tang, J., Jia, K.: Deep mesh reconstruction from single rgb images via topology modification networks. In: *ICCV* (2019) 4
51. Parkhi, O.M., Vedaldi, A., Zisserman, A., Jawahar, C.: Cats and dogs. In: *CVPR* (2012) 12
52. Paysan, P., Knothe, R., Amberg, B., Romdhani, S., Vetter, T.: A 3d face model for pose and illumination invariant face recognition. In: *2009 sixth IEEE international conference on advanced video and signal based surveillance*. pp. 296–301. *Ieee* (2009) 4, 7, 12
53. Pearl, J.: *Causality*. Cambridge university press (2009) 2
54. Phong, B.T.: Illumination for computer generated pictures. *Communications of the ACM* **18**(6), 311–317 (1975) 4
55. Schölkopf, B., Janzing, D., Peters, J., Sgouritsa, E., Zhang, K., Mooij, J.: On causal and anticausal learning. In: Langford, J., Pineau, J. (eds.) *Proceedings of the 29th International Conference on Machine Learning (ICML)*. pp. 1255–1262. Omnipress, New York, NY, USA (2012), <http://icml.cc/2012/papers/625.pdf> 6
56. Schölkopf, B., Locatello, F., Bauer, S., Ke, N.R., Kalchbrenner, N., Goyal, A., Bengio, Y.: Toward causal representation learning. *Proceedings of the IEEE* **109**(5), 612–634 (2021) 6
57. Schroff, F., Kalenichenko, D., Philbin, J.: Facenet: A unified embedding for face recognition and clustering. In: *CVPR* (2015)
58. Shen, X., Liu, F., Dong, H., Lian, Q., Chen, Z., Zhang, T.: Disentangled generative causal representation learning. *arXiv preprint arXiv:2010.02637* (2020) 7
59. Simonyan, K., Zisserman, A.: Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556* (2014)
60. Sinha, S.N., Ramnath, K., Szeliski, R.: Detecting and reconstructing 3d mirror symmetric objects. In: *ECCV* (2012) 4
61. Suwajanakorn, S., Snavely, N., Tompson, J.J., Norouzi, M.: Discovery of latent 3d keypoints via end-to-end geometric reasoning. *NeurIPS* (2018) 1
62. Tewari, A., Zollhofer, M., Kim, H., Garrido, P., Bernard, F., Perez, P., Theobalt, C.: Mofa: Model-based deep convolutional face autoencoder for unsupervised monocular reconstruction. In: *ICCV* (2017) 3
63. Thrun, S., Wegbreit, B.: Shape from symmetry. In: *ICCV* (2005) 4
64. Tulsiani, S., Zhou, T., Efros, A.A., Malik, J.: Multi-view supervision for single-view reconstruction via differentiable ray consistency. In: *CVPR* (2017) 4
65. Vowels, M.J., Camgoz, N.C., Bowden, R.: D’ya like dags? a survey on structure learning and causal discovery. *arXiv preprint arXiv:2103.02582* (2021) 6
66. Wang, C., Buenaposada, J.M., Zhu, R., Lucey, S.: Learning depth from monocular videos using direct methods. In: *CVPR* (2018) 4

67. Wang, M., Shu, Z., Cheng, S., Panagakis, Y., Samaras, D., Zafeiriou, S.: An adversarial neuro-tensorial approach for learning disentangled representations. *IJCV* **127**(6), 743–762 (2019) [4](#)
68. Wang, N., Zhang, Y., Li, Z., Fu, Y., Liu, W., Jiang, Y.G.: Pixel2mesh: Generating 3d mesh models from single rgb images. In: *ECCV* (2018) [4](#)
69. Wang, Q., et al.: Exponential convergence of the deep neural network approximation for analytic functions. *arXiv preprint arXiv:1807.00297* (2018)
70. Weichwald, S., Schölkopf, B., Ball, T., Grosse-Wentrup, M.: Causal and anti-causal learning in pattern recognition for neuroimaging. In: *International Workshop on Pattern Recognition in Neuroimaging* (2014) [6](#)
71. Wen, C., Zhang, Y., Li, Z., Fu, Y.: Pixel2mesh++: Multi-view 3d mesh generation via deformation. In: *ICCV* (2019) [4](#)
72. Wen, Y., Liu, W., Raj, B., Singh, R.: Self-supervised 3d face reconstruction via conditional estimation. In: *ICCV* (2021) [4](#)
73. Wiles, O., Zisserman, A.: Silnet: Single-and multi-view reconstruction by learning from silhouettes. In: *BMVC* (2017) [4](#)
74. Witkin, A.P.: Recovering surface shape and orientation from texture. *Artificial intelligence* **17**(1-3), 17–45 (1981) [4](#)
75. Wu, S., Rupprecht, C., Vedaldi, A.: Unsupervised learning of probably symmetric deformable 3d objects from images in the wild. In: *CVPR* (2020) [1](#), [3](#), [4](#), [8](#), [9](#), [10](#), [12](#), [13](#), [14](#)
76. Xiang, Y., Kim, W., Chen, W., Ji, J., Choy, C., Su, H., Mottaghi, R., Guibas, L., Savarese, S.: Objectnet3d: A large scale database for 3d object recognition. In: *ECCV* (2016) [4](#)
77. Xie, H., Yao, H., Sun, X., Zhou, S., Zhang, S.: Pix2vox: Context-aware 3d reconstruction from single and multi-view images. In: *ICCV* (2019) [4](#)
78. Yan, X., Yang, J., Yumer, E., Guo, Y., Lee, H.: Perspective transformer nets: Learning single-view 3d object reconstruction without 3d supervision. In: *NIPS* (2016) [4](#)
79. Yang, M., Liu, F., Chen, Z., Shen, X., Hao, J., Wang, J.: Causalvae: Disentangled representation learning via neural structural causal models. In: *CVPR* (2021) [7](#)
80. Yi, H., Li, C., Cao, Q., Shen, X., Li, S., Wang, G., Tai, Y.W.: Mmface: A multi-metric regression network for unconstrained face reconstruction. In: *CVPR* (2019) [4](#)
81. Yu, Y., Chen, J., Gao, T., Yu, M.: Dag-gnn: Dag structure learning with graph neural networks. In: *ICML* (2019) [11](#)
82. Zhang, R., Tsai, P.S., Cryer, J.E., Shah, M.: Shape-from-shading: a survey. *TPAMI* **21**(8), 690–706 (1999) [4](#)
83. Zhang, W., Sun, J., Tang, X.: Cat head detection-how to effectively exploit shape and texture features. In: *ECCV* (2008) [12](#)
84. Zheng, X., Aragam, B., Ravikumar, P.K., Xing, E.P.: Dags with no tears: Continuous optimization for structure learning. In: *NeurIPS* (2018) [11](#)
85. Zhou, T., Brown, M., Snavely, N., Lowe, D.G.: Unsupervised learning of depth and ego-motion from video. In: *CVPR* (2017) [4](#)
86. Zhu, R., Kiani Galoogahi, H., Wang, C., Lucey, S.: Rethinking reprojection: Closing the loop for pose-aware shape reconstruction from a single image. In: *ICCV* (2017) [4](#)
87. Zoph, B., Le, Q.V.: Neural architecture search with reinforcement learning. In: *ICLR* (2017) [11](#)