Leveraging Systematic Knowledge about 2D Transformations

Jiachen Kang¹ Wenjing Jia¹ Xiangjian He¹

¹School of Electrical and Data Engineering, University of Technology Sydney

Abstract

The existing deep learning models suffer from out-of-distribution (*o.o.d.*) performance drop in computer vision tasks. In comparison, humans have a remarkable ability to interpret images, even if the scenes in the images are rare, thanks to the systematicity of acquired knowledge. This work focuses on 1) the acquisition of systematic knowledge about 2D transformations, and 2) architectural components that can leverage the learned knowledge in image classification tasks. With a new training methodology based on synthetic datasets that are constructed under a causal framework, the deep neural networks acquire knowledge from semantically different domains (*e.g.* even from noise), and exhibit certain level of systematicity in parameter estimation experiments. Based on this, a novel architecture is devised consisting of a CLASSIFIER, an ESTIMATOR and an IDENTIFIER (abbreviated as "CED"). By emulating the "hypothesis-verification" process in human visual perception, CED improves the classification accuracy significantly on test sets under covariate shift.

1 Introduction

Machine learning algorithms based on deep neural networks (DNNs) have made dramatic progress in the field of computer vision in the last decade. Most of these algorithms strongly rely on the assumption of *i.i.d.*, *i.e.*, the training data and test data are independent and identically distributed. In practice, however, the *i.i.d.* assumption can be easily violated due to covariate shift in test datasets [1, 4, 15, 17], which causes significant performance drop of the models learned from the training set. This is investigated as the *o.o.d.* generalization problem, which has become one of the main challenges that the deep learning community encounters nowadays. One of the common stopgaps for this problem is to continuously expand the size of datasets, in order to strengthen the learned invariance of the target objects, by getting rid of other mechanisms or factors of variation. For example, ImageNet [10], which is a typical dataset for training classification and detection algorithms, contains more than 14 million images. Even so, popular classification models trained with ImageNet have



Figure 1: What is in image (a)? There are at least two ways to interpret it, *i.e.*, (b) three black circles partly *covered* by a white triangle, or (c) three black circles with a *notch* on each of them. (The former one may have a stronger tendency in perception, according to the Gestalt principles [19].)

experienced 40 - 45% performance drop when tested on ObjectNet, a biascontrolled dataset [4] that produces thousands of images with 600 combinations of parameters, by intervening only on three mechanisms in the photo generation process. This implies that if we try to construct a big enough dataset to approximate the distribution of real-world data, by considering all possible combinations of parameters of mechanisms, the number of required data points would be nearly infinite.

Human beings, in comparison, have powerful o.o.d. generalization abilities that enable us to recognize objects based on efficient learning. Extensive studies have shown that learned knowledge can be flexibly reused by infants in novel scenarios [39, 5, 22, 34]. This is analogous to algebraic operations [28] where symbolic variables are manipulated in computational processes. This can be a crucial explanation for the generalization ability. To illustrate this, if we look at Fig. 1(a) [24], at least two interpretations can be made (Fig. 1(b) and 1(c)), based on the same observation. This simple example illustrates a typical process of image perception, in which causal inference (in the anti-causal direction) is made by utilizing the mechanisms of either occlusion or notching on variables of circles and/or triangles. Specifically, the process consists of a hypothesis (of the content of three circles and a triangle) and the verification (whether a figure like this can be generated by *covering* the triangle over the circles). If another hypothesis (e.g. of just three circles) and a corresponding verification (by making a notch in each of them) can be made, the figure still makes sense to us. This "hypothesis-verification" process in human visual perception is discussed in detail in [27]. It can be noticed that mechanisms in the image generation process (e.g. occlusion or notching) are crucial in human visual perception. How an image is perceived relies on our knowledge of various mechanisms, rather than knowledge of images that are previously seen (which is the way that existing machines operate). It can also be noticed that our knowledge about occlusion or notching is universal and independent of the domain of variables. This generalization is also referred to as systematicity [12]. Based on the above analysis, it can be inferred that it is the systematicity of acquired knowledge that enables human beings to take mechanisms into consideration in visual perception, and thus achieve excellent o.o.d. generalization ability.

While children have plenty of time to gain systematic knowledge and phys-

ical mechanisms through observations and experiments [38, 35, 9], which build foundations for object perception and future knowledge acquisition [33, 37, 22], existing machine learning models rarely have opportunities to do so. One of the main reasons is that current datasets for visual learning inevitably introduce confounding mechanisms, which makes it difficult for models to learn unbiased representations and acquire systematic knowledge. Additionally, most of the studies focus on learning the invariance of objects of interest, and neglect the fact that other mechanisms also provide necessary information for perception, as shown in the previous example.

This work does not only focus on learning the invariance of objects of interest, we also pay attention to other mechanisms. Therefore, empirical studies are conducted to learn knowledge about mechanisms of 2D transformations (such as rotation, scaling and translation) using DNNs, in order to answer the following questions:

- 1. Whether the knowledge of these mechanisms learned by machines can exhibit some level of systematicity? If so,
- 2. whether the knowledge can be leveraged to facilitate image classification tasks like humans?

In order to answer the first question, it should be made clear what we mean by the knowledge of a mechanism. As human beings, for example, if we have learned the knowledge of 2D rotation, it means that for any image, (with a proper tool), (a) we can rotate the image at will, and (b) we are able to determine whether (and even how many degrees) the image has been rotated. Obviously, the knowledge we know about 2D rotation generalizes systematically and is independent of the domain of images. For transformations studied in this work, the affine transformation functions are in accord with the description in (a), and are used in the architecture as a tool to make precise operations¹. Therefore, our main purpose is the learning of the latter aspect (b). To achieve this, synthetic datasets are designed under a causal framework as the training datasets. Specifically, the datasets are composed of pairs of images, which are before and after the transformations, respectively. It has been found that with this training methodology, the transformation parameters can be estimated more accurately and stably even on *o.o.d.* data that are semantically different.

For the second research question, inspired by [27], the *hypothesis-verification* process in human perception is simulated in the task of hand-written digit classification, where modules of an ESTIMATOR and an IDENTIFIER trained offline separately, are used either as auxiliaries of the basic CLASSIFIER or as an independent architecture (Fig. 2). It is shown in the result that, by leveraging the learned knowledge of mechanisms, the ESTIMATOR and the IDENTIFIER as auxiliaries can improve the classification accuracy significantly with extra explainability. When the two modules operate independently of the CLASSIFIER,

¹It does not imply that transformation operations cannot be learned from data. Generative models which is beyond the scope of this study, have been studied in various tasks [29, 41].



Figure 2: The CED architecture. Potential classes are hypothesized by the CLASSIFIER C, and verification on these classes is made by the ESTIMATOR E and the IDENTIFIER D through the pipeline of (1) analyzing possible transformations, (2) reconstructing from candidates and (3) matching them with the sample.

without accessing any data of hand-written digits during training and through a pipeline of analyzing, reconstruction and matching, the architecture exceeds the performance of the basic CLASSIFIER.

To our best knowledge, this is the first work that utilize the systematic knowledge about other mechanisms in classifying images. As a result, in addition to answer questions like "Is there a '5' in the image? ", the proposed architecture is also able to answer "Why do you think it is a '5'? ", based on the knowledge it has been mastered, just like humans. The main contributions are as follows:

- We demonstrate a learning methodology, with which the DNNs can learn the knowledge of specific mechanisms robustly using synthetic datasets constructed under the causal framework (and thus the first research question is answered).
- We design a novel architecture that simulates human visual perception in image classification, with additional explainability, based on the knowledge that has been mastered (and answers the second question).

2 Related Work

In this section, techniques and research topics in computer vision related to this work, are briefly reviewed.

Data Augmentation and Domain Randomization. To tackle the potential drop in *o.o.d* performance, effective and commonly used techniques include data augmentation [14, 36] and domain randomization [40, 18, 32, 2]. These two techniques share similar principles. The former technique is usually referred specifically to 2D transformations; the latter is adopted when manipulations are made on parameters in 3D environments. In a causal perspective, they both make treatment randomization to get rid of confounders and to improve the learning of invariance. Based on this principle, this work also produces synthetic datasets through treatment randomization, but with a different purpose, that is, instead of randomizing *out* the mechanisms, we aim to take them *into* consideration in classification tasks.

Parameter Estimation. As introduced previously, the task for learning mechanisms of 2D transformations is to estimate the transformation parameters. This task is extensively studied in various computer vision topics, such as 2D spatial invariance learning [16], text detection [46, 45, 44, 43], and 3D pose estimations [26, 6], among many others. However, in most of the existing studies, parameter estimation can only be performed restricted to object categories that appear in the training sets. An important reason is that single-image parameter estimation is an ill-defined problem, in the sense that parameters in transformations are actually procedural variables, whose values are determined by both of the pre- and post-transformation states. Therefore, models trained with methodologies based on single images, can hardly generalize to unseen categories. In this work, the parameter estimation ability that we are interested in, should exhibit a certain degree of systematicity similar to human beings. Another series of works [48, 42] and the study in [25] conduct representation learning based on pairs of images that are related through mechanisms, by using a single encoder for multiple mechanisms. However we try to isolate knowledge about single mechanisms and reuse them in downstream tasks.

Program Induction. The knowledge learning in this work is essentially a program induction problem. Active deep learning topics in this area include program synthesis [3, 11], image generation [21, 47], *etc.* Program induction aims for more effective generation of programs, whereas this work focuses more on the interpretation of images that are beneficial for downstream tasks. Therefore, the domain-specific languages in this work are fundamentally different, being more semantically relevant to the downstream tasks.

3 Methodology

The aim of this work is to answer the two questions raised in the Introduction by investigating systematicity of the knowledge about the mechanisms and its application in hand-written digit classification. During classification, the test set has a potential covariate shift caused by a target mechanism that is known but cannot be overcome through data augmentation techniques (which is a common situation in real-world tasks). We simulate this setting by applying random 2D transformations on the MNIST [23] test set, with no data augmentation operations of any kind performed during training.

Inspired by the perception process in Fig. 1, we propose that if machines learn the knowledge of a target mechanism, they could perform better in classification under covariate shift that is caused by the mechanism. Hence, an architecture is devised consisting of three DNN modules: a CLASSIFIER C, an ESTIMATOR E and an IDENTIFIER D, and thus abbreviated as "CED". Causal



Figure 3: The causal graph of image transformation. X_t : Image at time step t (before transformation). X_{t+1} : Image at time step t+1 (after transformation). Θ : Parameter(s) of the transformation in study, as the variable is randomly sampled, this 'treatment randomization' operation removes all arrows pointing to Θ . U: Other unobservable variables that cause the generation of X_t .

datasets are constructed (in Section 3.1) for the modules E and D to learn the knowledge of mechanisms (in Section 3.2). CED makes predictions in classification by raising hypotheses with C and verifying them with E and D. The roles of these three modules are described in detail in Section 3.3.

3.1 Causal Datasets

To help DNNs learn the knowledge of a mechanism, the principle based on which a causal dataset is constructed, is explained below. Denote by \mathbf{x}_t and \mathbf{x}_{t+1} , respectively, the images before and after transformations f (parameterized with θ), *i.e.*,

$$\mathbf{x}_{t+1} = f(\mathbf{x}_t; \theta). \tag{1}$$

As explained in the Introduction, the goal of knowledge learning is to estimate the value of θ . Let X_t , X_{t+1} and Θ be the variables from which \mathbf{x}_t , \mathbf{x}_{t+1} and θ are instantiated, respectively. According to the causal graph in Fig. 3, if the estimation is made based only on the image *after* transformation, *i.e.*, $\mathbb{E}(\Theta|X_{t+1})$, given that X_{t+1} is a collider, conditioning on it will inevitably cause the information flow from U to Θ , which will hinder us from learning robust knowledge of f (via Θ). Therefore, in order to remove confounding caused by U, thus making the prediction of Θ more stable and generalize better in test domains, we have to condition on both X_t and X_{t+1} , *i.e.*, the Markov blanket of Θ .²

Concretely, in knowledge learning we aim to compute $\mathbb{E}_{P_{test}}(\Theta|X_t, X_{t+1})$ given only access to $P_{train}(\mathbf{x}_t, \mathbf{x}_{t+1}, \theta)$, with the assumptions of:

$$P_{train}(\theta | \mathbf{x}_t, \mathbf{x}_{t+1}) = P_{test}(\theta | \mathbf{x}_t, \mathbf{x}_{t+1}), \text{ and}$$
$$supp_{train}(\mathbf{x}_t, \mathbf{x}_{t+1}) = supp_{test}(\mathbf{x}_t, \mathbf{x}_{t+1}),$$

where P_{train} and P_{test} are distributions of data in training and test domains, respectively, and $P_{train}(\mathbf{x}_t, \mathbf{x}_{t+1}, \theta) \neq P_{test}(\mathbf{x}_t, \mathbf{x}_{t+1}, \theta)$.

In this work, synthetic datasets for knowledge learning are constructed according to the above causal framework. Each data point is composed of a pair

 $^{^{2}}$ This is also intuitively true, because it is pointless to ask how a picture has been transformed when no reference is provided.

of images that are before and after transformations and the label θ . Since the labels are automatically generated and no manual annotation is needed, this can be viewed as a self-supervised learning problem.

To acquire knowledge of mechanisms that is useful in classification with CED, one of the target mechanisms is the 2D transformation f_T that includes rotation, scaling and translation, and data points are generated through $\mathbf{x}_{t+1} = f_T(\mathbf{x}_t; \theta_T)$ with affine transformation functions as f_T . Another one is an identity function $f_I(\mathbf{x}_t; \theta_I)$ defined as:

$$\mathbf{x}_{t+1} = \begin{cases} f_T(\mathbf{x}_t; \hat{\theta}_T) & \text{if } \theta_I = 1; \\ f_T(\mathbf{x}_t'; \hat{\theta}_T) & \text{if } \theta_I = 0, \end{cases}$$
(2)

where \mathbf{x}'_t is a random sample other than \mathbf{x}_t , and $\hat{\theta}_T$ is the output of the module E in CED (which will be explained in detail in Sections 3.2 and 3.3).

3.2 Knowledge Learning

Based on the above causal datasets, the ESTIMATOR E and the IDENTIFIER D are trained to learn knowledge of 2D transformations f_T and the identity function f_I , respectively. Specifically, we employ E that takes as the input paired images \mathbf{x}_t^E and \mathbf{x}_{t+1}^E generated from f_T to predict the parameters $\hat{\theta}_T$. The role of D, on the other hand, is to learn from f_I and to predict the probability that two images are of the same identity. In practice, the inputs of D are $\mathbf{x}_t^D = \mathbf{x}_{t+1}^E = f_T(\mathbf{x}_t^E; \theta_T)$, and $\mathbf{x}_{t+1}^D = f_I(\mathbf{x}_t^E; \theta_I)$. The mechanism of f_T are independent of f_I , and thus E is optimized first,

The mechanism of f_T are independent of f_I , and thus E is optimized first, by minimizing the mean squared error (L_{MSE}) . D is then trained based on datasets generated with f_I and E, and optimized by minimizing the binary cross entropy loss (L_{BCE}) . Therefore, the objectives of knowledge learning in this study can be represented as:

$$\arg\min_{E} L_{MSE}(E(\mathbf{x}_{t}^{E}, f_{T}(\mathbf{x}_{t}^{E}; \theta_{T})), \theta_{T})$$
(3)

$$\arg\min_{D} L_{BCE}(D(f_T(\mathbf{x}_t^E; \theta_T), f_I(\mathbf{x}_t^E; \theta_I)), \theta_I)$$
(4)

Models To obtain a more robust DNN in knowledge learning for modules E and D, three Convolutional Neural Networks (CNNs) are investigated, as illustrated in Fig. 4. The first model in study takes concatenated images as input (shown in Fig. 4(a)). It is called FactorNet for brevity in this work. The following two models are relevant to this work, and thus are used as baselines. The Siamese Networks [7] (shown in Fig. 4(b)) are extensively studied on datasets with intrinsic relations in metric learning and representation learning. Vanilla CNN (shown in Fig. 4(c)) that takes single-images as input to make predictions, is another common method for numerical regression tasks.



Figure 4: The models in knowledge learning. (a) FactorNet: \mathbf{x}_t and \mathbf{x}_{t+1} are concatenated in channel dimension before being fed into CNN; (b) Siamese Network: \mathbf{x}_t and \mathbf{x}_{t+1} are fed into CNN, whose outputs are then concatenated and sent to fully-connected (FC) layers; (c) Vanilla CNN: Only the transformed images \mathbf{x}_{t+1} are fed into CNN.

3.3 Architecture CED for Classification

We now describe in detail each module in the proposed architecture CED for classification and their roles in simulating the "hypothesis-verification" process.

Classifier C. The images in the MNIST test set are transformed before testing, denoted by X_t^{test} , while those in the training set are original ones without any transformation, denoted by X_0^{train} . Given a test sample $\mathbf{x}_t^{test} \in X_t^{test}$, module C produces a probability distribution across all classes, which is exploited as confidence scores. If the highest confidence score is lower than a preset threshold, instead of making a prediction of label, C will output a hypothesis $H(\mathbf{x}_t^{test}) = \{y_i\}_{i=1}^k$, containing a list of class labels with top k(k > 1) confidence scores.

Estimator *E*. Module *E* randomly samples $N(N \ge 1)$ candidates from X_0^{train} for each class in $H(\mathbf{x}_t^{test})$. Concretely, if the set of all candidates for \mathbf{x}_t^{test} is denoted by $X_c \subset X_0^{train}$, we have $X_c = \{X_c^{(y_i)}\}_{i=1}^k$, and each $X_c^{(y_i)} = \{\mathbf{x}_j^{(y_i)}\}_{j=1}^N$. With the assumption that \mathbf{x}_t^{test} may be transformed from what looks similar to some of the candidates in X_c , *E* then analyzes the relationship between \mathbf{x}_t^{test} and each candidate *w.r.t.* 2D transformation using knowledge

learned previously, by computing $\widehat{\theta}_T^{i,j} = E(\mathbf{x}_j^{(y_i)}, \mathbf{x}_t^{test})$. **Identifier** *D*. Since *E* is a deterministic function and will produce an out-

Identifier *D*. Since *E* is a deterministic function and will produce an output, regardless of whether two images are really related, the role of *D* is to examine which candidate is more closely related to the \mathbf{x}_t^{test} . To achieve this, firstly, *D* performs reconstructions on each candidate by exploiting the instructions from *E* and obtains $\mathbf{x}_{j,t}^{(y_i)} = f_T(\mathbf{x}_j^{(y_i)}; \widehat{\theta}_T^{i,j})$. Then, $\mathbf{x}_{j,t}^{(y_i)}$ is tested on how likely it matches to \mathbf{x}_t^{test} , by leveraging the knowledge learned about the identity function f_I . The label of the candidate with highest likelihood will be output as the final prediction $\hat{y} = \arg \max_{y_i} D(\mathbf{x}_t^{test}, \mathbf{x}_{j,t}^{(y_i)})$.

In the above process, potential classes are hypothesized by C, and verification on these classes is made by modules E and D through the pipeline of (a) analyzing possible transformations, (b) reconstructing from candidates and (c) matching them with the sample.

ED. It can also be noticed that the pre-trained modules E and D do not have to access MNIST during training, and do not rely on C too much either. Based on the fact that the training and test set of MNIST share the same class label space, we also explore a second architecture that only employs E and D (abbreviated as "ED"). The only difference from CED is that ED directly takes all classes as hypothesis (k = 10).

4 Experiments

In this section, experiments are conducted to answer the two questions raised in the Introduction.

4.1 Is the Learned Knowledge Systematic?

In order to study the robustness of estimation on θ_T and θ_I of f_T and f_I , synthetic datasets are constructed according to Section 3.1. Three DNN models are trained and tested based on the methodology illustrated in Section 3.2. We now describe specifically how the experiments are conducted.

4.1.1 Learning of 2D Image Transformation mechanisms

Datasets. In the experiments, the original images in MNIST, EMNIST [8] and CIFAR-10 [20] are used as \mathbf{x}_0 . To eliminate potential overfitting, we obtain the input image pairs $\mathbf{x}_t = f_T(\mathbf{x}_0; \theta_T^0)$ and $\mathbf{x}_{t+1} = f_T(\mathbf{x}_t; \theta_T)$, where θ_T^0 and θ_T are randomly sampled in a uniform distribution (see Table 1).

In this work, we conduct learning on four types of f_T , including the individual learning of rotation, scaling and translation, and the joint learning of all the above three. For individual learning, only one of the three transformations is applied on \mathbf{x}_t at a time, while all three transformations are applied simultaneously for joint learning.

To further increase the difficulty of the task, a synthetic dataset composed of black/white noises (of a Bernoulli distribution) is randomly generated and

Parameter	Range
Rotation angle	$[-90^{\circ}, 90^{\circ}]$
Translation (horizontal)	[-5, 5] pixels
Translation (vertical)	[-5, 5] pixels
Scale factor	[0.7, 1.3]

Table 1: The parameters of 2D transformations. The values of each parameter are uniformly sampled within their ranges.

Table 2: The training and test data used for knowledge learning.

Experiment		Dataset
Exp_MNIST	Train	training set in MNIST
	Test	'letter' division of test
		set in EMNIST
Exp_CIFAR	Train	9 classes of training set
		in CIFAR-10
	Test	the other class of
		training set in CIFAR-10
Exp_NOISE	Train	black/white noise
	Test	test set in MNIST

used as \mathbf{x}_t . To better test robustness, all test data are sampled from datasets that are semantically different from the training sets. The detailed schemes are listed in Table 2.

Models. The CNN backbone in [48] is used in all three models in Fig. 4. All input pairs of \mathbf{x}_t and \mathbf{x}_{t+1} are concatenated along the channel dimension before being fed in to the FactorNets; the input size is $N_{batch} \times 2 \times 28 \times 28$ in Exp_MNIST and Exp_NOISE, and $N_{batch} \times 6 \times 32 \times 32$ in Exp_CIFAR, where N_{batch} is the batch size.

The performance of FactorNets for 2D transformations learning is reported in Fig. 5 and Fig. 9 in Appendix A. For the learning of individual mechanisms, it can be observed that most of the absolute percentage errors (APE) (*e.g.* the third quartile in the distributions) can be controlled to below 20% in most of the experiments, and even 10% in Exp_CIFAR. The accurate parameter estimation suggests effective learning of 2D transformation knowledge. Furthermore, there are only minor differences between the distributions of APE when we compare the training and test sets. This can be attributed to the strong *o.o.d.* systematicity of knowledge learned with FactorNets, given that the data in the training and test sets are completely different in semantics. More results on the performance of FactorNets for 2D transformation learning are shown in Appendix A.



Figure 5: Performance of FactorNets for individual rotation learning. (left) Predictions of rotation angle vs. the ground truth (normalized to [-1, 1]) in test set. (right) Distributions of absolute percentage errors (in %) of all data points in the dataset.

4.1.2 Learning of the Identity Function

To evaluate the *o.o.d.* performance of the θ_I estimation, we train FactorNet in Exp_NOISE using samples generated with the method in Section 3.1, and θ_I is randomly sampled in $\{0, 1\}$ in order to produce a balanced dataset. The FactorNet for rotation learning trained in Exp_NOISE is used as E. The resulting F1 scores are 0.9987 and 0.9757 for training and test set, respectively, which indicates superior *o.o.d.* performance of FactorNet for identification tasks as the module D in CED.

4.1.3 Key Elements in Knowledge Learning

In this section, several ablation studies are conducted to examine elements crucial for a robust knowledge learning.

Firstly, as analyzed based on the causal graph in Fig. 3, if there exists causal relationship from U to X_t , it is necessary to condition on both X_t and X_{t+1} in order to predict Θ robustly. As shown in Fig. 6, the *o.o.d.* performance gap of vanilla CNN is noticeable in all learning cases, compared with FactorNet and Siamese networks that take both X_t and X_{t+1} as inputs. Vanilla CNN performs relatively better in translation learning, because the position of X_t (the original images in this case) is always in the center and independent of U. However, while being able to estimate rotation angles accurately in the training set, vanilla CNN completely fails in the test set, because the estimation of angles relies highly on the pattern of images, which is determined by U. This also offers insight into numerical regression tasks in contemporary computer vision studies, such as object pose estimation; given only the images after transformation during training, a good *o.o.d.* performance cannot be expected.

Secondly, for CNN backbones, computations based on concatenated images are necessary to make more accurate estimations. Fig. 6 shows that Siamese networks underperform FactorNets in all mechanisms. Much information about transformations is lost through convolutional operations and the max pooling



Figure 6: The performance of individual transformation learning across different models.

layers, while more information can be preserved from the beginning in Factor-Nets.

Additionally, we speculate that the inductive bias of CNNs fundamentally affects the effectiveness of knowledge learning. This is based on the observation of the learning curves of the three mechanisms (in Fig. 12 in Appendix A). Fast learning on translation and scaling and a slow one on rotation can be noticed for all models, which indicates that CNN models have greater difficulty learning the mechanism of rotation. Considering CNN properties of translation-equivariant, positional information can be encoded and operated with CNN at higher efficiency. An extensive investigation into other inductive bias is necessary for a more solid claim to be made in the future.

4.2 Can Knowledge be Leveraged?

In the previous section, it can be seen that effective learning can be achieved with FactorNets. The models are capable of making accurate estimations on parameters θ , and this capability can be generalized to semantically different datasets. This indicates a certain level of systematicity. Hence, with these models as building blocks, we construct the CED architecture according to Section 3.3. Comparison of classification performance amongst different architectures are conducted first (in Section 4.2.1), followed by discussion of the simulation of human-like visual perception (in Section 4.2.2).

4.2.1 Classification Performance

In the experiment, classification is performed with the setting of covariate shift caused by rotation, and the performance is compared amongst a basic classifier, CED and ED. To construct CED and ED, FactorNets trained in Exp_NOISE for the (individual) rotation learning and the identity function learning are exploited as the module E and D, respectively. The basic classifier C is trained with original images X_0^{train} in MNIST without any data augmentations. The length of hypothesis $H(\mathbf{x}_t^{test})$ is k = 5, 10 for CED and k = 10 for ED. The



Figure 7: The performance of classification. CED_5 and CED_10 denote CED with hypothesis k = 5 and k = 10, respectively.

number of candidates for the E is N = 200 for each class. The confidence threshold of C is set to 0.9999.

The classification accuracy obtained on the MNIST test set, with or without rotations, is shown in Fig. 7. The first observation is that, in the case of rotated test set, the basic classifier has experienced nearly a 40% performance drop. However, the accuracy of CED has increased to 77% when k = 5 (CED_5) and further to 82% when k = 10 (CED_10). In CED, E and D are introduced for further interpretation when C is not very confident, and they provide extra explanations about why the sample is classified as such and how it is rotated, by leveraging the knowledge about rotation with E. Additionally, this process does not affect the performance too much for the test set without rotation.

ED outperforms the basic classifier by classifying with an accuracy of more than 75%. It is worth noting that the performance is achieved without any knowledge of the handwritten digits (since both E and D are trained in Exp_NOISE), but only through the processes of analyzing, reconstructing and matching. Furthermore, only 4% ($200 \times 10/50000$) of the training data are accessed during inference. This is behaviorally similar to human beings, who are capable of classifying characters that they do not know at all, so long as necessary references are provided.

To investigate further the role of E with its knowledge about rotation, an ablation study was conducted on CD by removing E from the CED. Obviously, the CD architecture loses the ability to interpret transformation and the performance on rotated test set has dropped to below 60% (Fig. 7). On one hand, this indicates the importance of rotation knowledge to D, which requires the instructions for reconstruction; On the other hand, since the rotated samples look very different from the candidates, indirectly it also demonstrates the effectiveness of D.

The number of candidates. As shown in Fig. 8, classification accuracy is greatly affected by the number of candidates. Given that D is trained on



Figure 8: The classification accuracy of CED with different numbers of candidates. Performance already surpasses the basic classifier (the green dash line) when $N \ge 10$.

noise, the module is really sensitive to nuance differences. Therefore, in order to find a candidate that is very similar to a sample, a very large candidate pool is required.

In addition, the generation of digits can also be viewed as a mechanism. Unlike 2D transformations, the parameterization of digit generation is much more complicated [21]. While the integration of an estimation module for digit generation (as a new E) into the existing CED would presumably reduce the required number of candidates significantly, this will, at the same time, introduce new challenges in compositionality, which involves the collaboration between multiple Es.

4.2.2 Simulation of Human-like Visual Perception

In this work, we propose CED as a preliminary simulation of the "hypothesisverification" process [27] in human visual perception. Although the simulation is not a reverse engineering of the human brain, based on psychological studies about cognition and behaviors, both human and CED share similarities in how information is processed.

As human beings, we have powerful ability to model an object with functionally easier mechanisms according to Gestalt principles [19]. This does not happen only in visual perception, but also in other aspects of behaviors [13, 30], where people try to rationalize their behaviors with convincing (but sometimes incorrect) reasons. The role of E and D in CED is actually to provide explainability, with which machines can "make sense", to some extent, of what they see. This explainability also provides possibilities for humans to improve the architectures, in ways that they can comprehend.

Furthermore, the simulation and imagination in the brain have been studied

in various works, and are proposed as the key elements in the understanding of physical scenes and counterfactual reasoning [5, 31]. Based on the model of the world in the mind, humans can make predictions about the future (in causal direction) and infer the causes of things that have happened (in anti-causal direction). In the architecture of CED, simulations of 2D-transformations in anti-causal and causal directions are enabled with module E and the affine transformation function, respectively, which equip the machine with an imagination space.

5 Conclusion and Future Work

To conclude, in this study, stable knowledge learning has been shown to be possible if models have been trained on (concatenated) data pairs that are intrinsically related through the mechanism. Based on this learning methodology, FactorNets with their acquired knowledge play significant roles in image classifications and further interpretations under covariate shift. The performance boost of the proposed CED architecture also suggests the effectiveness of the simulation to human-like visual perception. We hope the simulation, along with its basis, *i.e.*, the learning methodology, can provide inspirations for future studies in computer vision w.r.t. human-like general AI.

Based on the findings in this work, we identify some limitations and questions that are appealing for further investigations.

Compositionality: In this work, covariate shift is introduced in test set by intervening on only one mechanism (*i.e.* rotation). In the setting where multiple mechanisms are considered, it will be ideal if multiple Es could leverage the knowledge learned separately and cooperate with each other. However, some preliminary results show that Es will not generalize well, if the training is based on only the interventions of target mechanism and keeping the others fixed. This is in line with [26], where the generalization improves only if more combinations of two mechanisms (category and pose) are exposed *during training*. Therefore, essential architectural elements that would facilitate the communications and interactions between modules (especially Es) are intriguing for us to explore in the future.

3D Virtual World: If we think of the real-world photos as the result of the interactions of mechanisms, such as foreground and background objects, lighting conditions, camera attributes, *etc.*, then tasks based on real photos could also be tackled in the same manner as in this work. With the rapid development of computer graphics, photo-realistic synthetic datasets with 1) controlled interventions on target mechanisms and 2) automatic pixel-accurate annotations can be efficiently created with 3D rendering engines. As described in Section 3.1, if the mechanism is stable across both virtual and real worlds, the knowledge learned on synthetic images could presumably be usable on real photos. More importantly, the real-time rendering capability of modern game engines (*e.g.* Unreal Engine) offers potential realization of an imagination space for machines (analogous to the affine transformation functions in the work).

References

- [1] Michael A Alcorn, Qi Li, Zhitao Gong, Chengfei Wang, Long Mai, Wei-Shinn Ku, and Anh Nguyen. Strike (with) a pose: Neural networks are easily fooled by strange poses of familiar objects. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4845–4854, 2019.
- [2] Artemij Amiranashvili, Max Argus, Lukas Hermann, Wolfram Burgard, and Thomas Brox. Pre-training of deep rl agents for improved learning under domain randomization. arXiv preprint arXiv:2104.14386, 2021.
- [3] Matej Balog, Alexander L. Gaunt, Marc Brockschmidt, Sebastian Nowozin, and Daniel Tarlow. Deepcoder: Learning to write programs. In 5th International Conference on Learning Representations, ICLR 2017, Toulon, France, April 24-26, 2017, Conference Track Proceedings. OpenReview.net, 2017.
- [4] Andrei Barbu, David Mayo, Julian Alverio, William Luo, Christopher Wang, Dan Gutfreund, Josh Tenenbaum, and Boris Katz. Objectnet: A large-scale bias-controlled dataset for pushing the limits of object recognition models. In Advances in Neural Information Processing Systems, volume 32. Curran Associates, Inc., 2019.
- [5] Peter W Battaglia, Jessica B Hamrick, and Joshua B Tenenbaum. Simulation as an engine of physical scene understanding. *Proceedings of the National Academy of Sciences*, 110(45):18327–18332, 2013.
- [6] Tsai-Shien Chen, Man-Yu Lee, Chih-Ting Liu, and Shao-Yi Chien. Aware channel-wise attentive network for vehicle re-identification. In *Proceedings* of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops, pages 574–575, 2020.
- [7] Sumit Chopra, Raia Hadsell, and Yann LeCun. Learning a similarity metric discriminatively, with application to face verification. In 2005 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR'05), volume 1, pages 539–546. IEEE, 2005.
- [8] Gregory Cohen, Saeed Afshar, Jonathan Tapson, and Andre Van Schaik. Emnist: Extending mnist to handwritten letters. In 2017 International Joint Conference on Neural Networks (IJCNN), pages 2921–2926. IEEE, 2017.
- [9] Claire Cook, Noah D Goodman, and Laura E Schulz. Where science starts: Spontaneous experiments in preschoolers' exploratory play. *Cognition*, 120(3):341–349, 2011.
- [10] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In 2009 IEEE conference on computer vision and pattern recognition, pages 248–255. IEEE, 2009.

- [11] Kevin Ellis, Catherine Wong, Maxwell Nye, Mathias Sablé-Meyer, Lucas Morales, Luke Hewitt, Luc Cary, Armando Solar-Lezama, and Joshua B. Tenenbaum. DreamCoder: Bootstrapping Inductive Program Synthesis with Wake-Sleep Library Learning, page 835–850. Association for Computing Machinery, New York, NY, USA, 2021.
- [12] Jerry A Fodor and Zenon W Pylyshyn. Connectionism and cognitive architecture: A critical analysis. *Cognition*, 28(1-2):3–71, 1988.
- [13] Michael S Gazzaniga. The split brain revisited. Scientific American, 279(1):50–55, 1998.
- [14] Ian Goodfellow, Yoshua Bengio, and Aaron Courville. Deep learning. MIT press, 2016.
- [15] Ian J Goodfellow, Jonathon Shlens, and Christian Szegedy. Explaining and harnessing adversarial examples. arXiv preprint arXiv:1412.6572, 2014.
- [16] Max Jaderberg, Karen Simonyan, Andrew Zisserman, et al. Spatial transformer networks. Advances in neural information processing systems, 28:2017–2025, 2015.
- [17] Jason Jo and Yoshua Bengio. Measuring the tendency of cnns to learn surface statistical regularities. arXiv preprint arXiv:1711.11561, 2017.
- [18] Rawal Khirodkar, Donghyun Yoo, and Kris Kitani. Domain randomization for scene-specific car detection and pose estimation. In 2019 IEEE Winter Conference on Applications of Computer Vision (WACV), pages 1932–1940. IEEE, 2019.
- [19] Kurt Koffka. Principles of Gestalt psychology. Routledge, 2013.
- [20] Alex Krizhevsky, Geoffrey Hinton, et al. Learning multiple layers of features from tiny images. Technical report, University of Toronto, 2009.
- [21] Brenden M Lake, Ruslan Salakhutdinov, and Joshua B Tenenbaum. Human-level concept learning through probabilistic program induction. *Science*, 350(6266):1332–1338, 2015.
- [22] Brenden M Lake, Tomer D Ullman, Joshua B Tenenbaum, and Samuel J Gershman. Building machines that learn and think like people. *Behavioral* and brain sciences, 40, 2017.
- [23] Yann LeCun, Léon Bottou, Yoshua Bengio, and Patrick Haffner. Gradientbased learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, 1998.
- [24] Steven M Lehar. The world in your head: A gestalt view of the mechanism of conscious experience. Psychology Press, 2003.

- [25] Francesco Locatello, Ben Poole, Gunnar Rätsch, Bernhard Schölkopf, Olivier Bachem, and Michael Tschannen. Weakly-supervised disentanglement without compromises. In *International Conference on Machine Learning*, pages 6348–6359. PMLR, 2020.
- [26] Spandan Madan, Timothy Henry, Jamell Dozier, Helen Ho, Nishchal Bhandari, Tomotake Sasaki, Frédo Durand, Hanspeter Pfister, and Xavier Boix. On the capability of neural networks to generalize to unseen category-pose combinations. Technical report, Center for Brains, Minds and Machines (CBMM), 2020.
- [27] Anthony J Marcel. Conscious and unconscious perception: An approach to the relations between phenomenal experience and perceptual processes. *Cognitive psychology*, 15(2):238–300, 1983.
- [28] Gary F Marcus. The algebraic mind: Integrating connectionism and cognitive science. MIT press, 2003.
- [29] Mehdi Mirza and Simon Osindero. Conditional generative adversarial nets. arXiv preprint arXiv:1411.1784, 2014.
- [30] Richard E Nisbett and Timothy D Wilson. Telling more than we can know: Verbal reports on mental processes. *Psychological review*, 84(3):231, 1977.
- [31] Judea Pearl and Dana Mackenzie. The book of why: the new science of cause and effect. Basic books, 2018.
- [32] Xinyi Ren, Jianlan Luo, Eugen Solowjow, Juan Aparicio Ojea, Abhishek Gupta, Aviv Tamar, and Pieter Abbeel. Domain randomization for active pose estimation. In 2019 International Conference on Robotics and Automation (ICRA), pages 7228–7234. IEEE, 2019.
- [33] Hilary Schmidt and Elizabeth Spelke. The development of gestalt perception in infancy. *Infant Behavior and Development*, 9:329, 1986.
- [34] Bernhard Schölkopf, Francesco Locatello, Stefan Bauer, Nan Rosemary Ke, Nal Kalchbrenner, Anirudh Goyal, and Yoshua Bengio. Toward causal representation learning. *Proceedings of the IEEE*, 109(5):612–634, 2021.
- [35] Laura E Schulz, Alison Gopnik, and Clark Glymour. Preschool children learn about causal structure from conditional interventions. *Developmental* science, 10(3):322–332, 2007.
- [36] Connor Shorten and Taghi M Khoshgoftaar. A survey on image data augmentation for deep learning. *Journal of Big Data*, 6(1):1–48, 2019.
- [37] Elizabeth S Spelke. Principles of object perception. Cognitive science, 14(1):29–56, 1990.
- [38] Aimee E Stahl and Lisa Feigenson. Observing the unexpected enhances infants' learning and exploration. *Science*, 348(6230):91–94, 2015.

- [39] Ernő Téglás, Edward Vul, Vittorio Girotto, Michel Gonzalez, Joshua B Tenenbaum, and Luca L Bonatti. Pure reasoning in 12-month-old infants as probabilistic inference. *science*, 332(6033):1054–1059, 2011.
- [40] Josh Tobin, Rachel Fong, Alex Ray, Jonas Schneider, Wojciech Zaremba, and Pieter Abbeel. Domain randomization for transferring deep neural networks from simulation to the real world. In 2017 IEEE/RSJ international conference on intelligent robots and systems (IROS), pages 23–30. IEEE, 2017.
- [41] Pascal Vincent, Hugo Larochelle, Isabelle Lajoie, Yoshua Bengio, Pierre-Antoine Manzagol, and Léon Bottou. Stacked denoising autoencoders: Learning useful representations in a deep network with a local denoising criterion. Journal of machine learning research, 11(12), 2010.
- [42] Xiao Wang, Daisuke Kihara, Jiebo Luo, and Guo-Jun Qi. Enaet: Selftrained ensemble autoencoding transformations for semi-supervised learning. arXiv preprint arXiv:1911.09265, 2, 2019.
- [43] Chengpei Xu, Hao Fu, Long Ma, Wenjing Jia, Chengqi Zhang, Feng Xia, Xiaoyu Ai, Binghao Li, and Wenjie Zhang. Seeing text in the dark: Algorithm and benchmark. arXiv preprint arXiv:2404.08965, 2024.
- [44] Chengpei Xu, Wenjing Jia, Tingcheng Cui, Ruomei Wang, Yuan-fang Zhang, and Xiangjian He. Arbitrary-shape scene text detection via visualrelational rectification and contour approximation. *IEEE Trans. Multimedia*, 2022.
- [45] Chengpei Xu, Wenjing Jia, Ruomei Wang, Xiaonan Luo, and Xiangjian He. Morphtext: Deep morphology regularized accurate arbitrary-shape scene text detection. *IEEE Trans. Multimedia*, 2022.
- [46] Chengpei Xu, Ruomei Wang, Shujin Lin, Xiaonan Luo, Baoquan Zhao, Lijie Shao, and Mengqiu Hu. Lecture2note: Automatic generation of lecture notes from slide-based educational videos. In 2019 IEEE International Conference on Multimedia and Expo (ICME), pages 898–903. IEEE, 2019.
- [47] Halley Young, Osbert Bastani, and Mayur Naik. Learning neurosymbolic generative models via program synthesis. In *International Conference on Machine Learning*, pages 7144–7153. PMLR, 2019.
- [48] Liheng Zhang, Guo-Jun Qi, Liqiang Wang, and Jiebo Luo. Aet vs. aed: Unsupervised representation learning by auto-encoding transformations rather than data. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pages 2547–2555, 2019.

A Additional Results

Individual learning. Additional results of performance of FactorNets for individual 2D transformation learning is shown in Fig. 9. Similar to the result in Fig. 5, several observations for individual learning are listed as follows.

- Majority of absolute percentage errors can be controlled to below 20% for individual learning, which indicates the effectiveness of 2D transformation learning.
- There are only minor differences in the distributions of absolute percentage error between the training and test sets for individual learning across all experiments, which suggests strong *o.o.d.* generalization.



Figure 9: Performance of FactorNets for individual 2D transformation learning. (left) Rotation. (center) Scaling. (right) Translation.

Joint learning. For joint learning of 2D transformation, obvious performance drop in both the training and test set can be observed in Fig. 10, compared with the individual learning, even if the number of parameters of FactorNets is four times that of models for individual learning. Similar results are reported in study [26], where more accurate estimations of variables are made by separately trained models, because of the improved "selectivity and invariance at the individual neuronal level".

FactorNets trained in Exp_NOISE. Although FactorNet exhibits strong o.o.d. generalization, the performance decreases to some extent when the difference between the training and test sets becomes considerably big. For instance, a larger performance gap between the training and test set in Exp_NOISE can be noticed, compared with the other two experiments in Fig. 9 and 10. The most apparent characteristic in this experiment is the pattern difference between noises and hand-written digits, which implies the potential difference in exploitation of patterns during learning.



Figure 10: Performance of FactorNets for joint 2D transformation learning. (left) Rotation. (center) Scaling. (right) Translation.

To prove this, an ablation study was conducted by altering the ratio of black to white pixels of the training data in Exp_NOISE. As shown in Fig. 11, the bestperforming model for rotation learning is trained on 7:3 black/white noises. However, if the pixel values in MNIST are swapped (*i.e.* black digits on white background), the best performance can be achieved around 4:6. Different ratios will provide different patterns that can be exploited in learning. The best ratio for individual learning of translation and rotation is around 7:3, while for scaling it is around 3:7, which can also explain the poor *o.o.d.* generalization performance of joint learning in Exp_NOISE, since it is impossible for the model to learn the three transformations equally well with only one ratio.



Figure 11: Performance of FactorNets in rotation learning with controlled black/white pixel ratios in EXP_NOISE. Pixel values are swapped in MNIST_b.

Learning curves in 2D transformation learning. The learning curves in 2D transformation learning are shown in Fig. 12. For all three models, fast learning on translation and scaling and a slow one on rotation can be observed.



Figure 12: The learning curves in transformation learning across different models. Fast learning on translation and scaling and a slow one on rotation can be observed for all models.

Restoration. An interesting property of FactorNet and Siamese networks can be found further in translation learning. Given an image \mathbf{x}_t with a small square in the center, an image \mathbf{x}_{t+1} identical to \mathbf{x}_t and the target value of translation θ , we can obtain a (coarse) translated version of \mathbf{x}_t by optimizing \mathbf{x}_{t+1} with gradient decent according to:

$$\mathbf{x}_{t+1} \leftarrow \mathbf{x}_{t+1} - \alpha \nabla_{\mathbf{x}_{t+1}} L_{MSE}(E(\mathbf{x}_t, \mathbf{x}_{t+1}), \theta), \tag{5}$$

where α is the learning rate. As shown in Fig. 13, this operation can be viewed as an approximation of the translation function f_T . Although this reversed generation of images is by no means accurate and only limited to very simple patterns, the phenomenon cannot be repeated in cases of rotation and scaling.

B Model Architecture Details

We follow the implementation in [48] to construct the three models (in Fig 4) for knowledge learning experiments. The architectures for individual mechanism learning are shown in Table 3. The models for joint learning are only different in channel sizes, which are all doubled in Exp_MNIST and Exp_NOISE, and 50% larger in Exp_CIFAR.



Figure 13: Images obtained with the Translation FactorNet through gradient decent. The image in the center is the original one \mathbf{x}_t . According to the values of θ (four of them are marked on the corners), \mathbf{x}_{t+1} are generated through gradient decent. In each of \mathbf{x}_{t+1} , an obvious offset of the light area from the original position (the blue dot) to the target position can be observed.

Table 3: Architecture of models for knowledge learning.

Models in Exp_MNIST and Exp_NOISE	Models in Exp_CIFAR
5×5 Conv 96, BatchNorm, ReLU	5×5 Conv 192, BatchNorm, ReLU
1×1 Conv 64, BatchNorm, ReLU	1×1 Conv 128, BatchNorm, ReLU
1×1 Conv 32, BatchNorm, ReLU	1×1 Conv 64, BatchNorm, ReLU
3×3 MaxPooling stride 2	3×3 MaxPooling stride 2
3×3 Conv 32, BatchNorm, ReLU	3×3 Conv 128, BatchNorm, ReLU
1×1 Conv 32, BatchNorm, ReLU	1×1 Conv 128, BatchNorm, ReLU
1×1 Conv 32, BatchNorm, ReLU	1×1 Conv 128, BatchNorm, ReLU
3×3 MaxPooling stride 2	3×3 MaxPooling stride 2
3×3 Conv 32, BatchNorm, ReLU	3×3 Conv 128, BatchNorm, ReLU
1×1 Conv 32, BatchNorm, ReLU	1×1 Conv 128, BatchNorm, ReLU
1×1 Conv 32, BatchNorm, ReLU	1×1 Conv 128, BatchNorm, ReLU
3×3 MaxPooling stride 2	3×3 MaxPooling stride 2
2×2 Conv 32, BatchNorm, ReLU	2×2 Conv 128, BatchNorm, ReLU
1×1 Conv 32, BatchNorm, ReLU	1×1 Conv 128, BatchNorm, ReLU
1×1 Conv 32, BatchNorm, ReLU	1×1 Conv 128, BatchNorm, ReLU
3×3 MaxPooling stride 2	3×3 MaxPooling stride 2
FC	FC
FC (Siamese Networks only)	FC (Siamese Networks only)