Feedback Networks

Amir R. Zamir[1,3*]   Te-Lin Wu[1*]   Lin Sun[1,2]   William B. Shen[1]   Bertram E. Shi[2]

Jitendra Malik[3]   Silvio Savarese[1]

[1] Stanford University   [2] HKUST   [3] University of California, Berkeley

http://feedbacknet.stanford.edu/

Abstract

Currently, the most successful learning models in computer vision are based on learning successive representations followed by a decision layer. This is usually actualized through feedforward multilayer neural networks, e.g. ConvNets, where each layer forms one of such successive representations. However, an alternative that can achieve the same goal is a feedback based approach in which the representation is formed in an iterative manner based on a feedback received from previous iteration's output.

We establish that a feedback based approach has several core advantages over feedforward: it enables making early predictions at the query time, its output naturally conforms to a hierarchical structure in the label space (e.g. a taxonomy), and it provides a new basis for Curriculum Learning. We observe that feedback develops a considerably different representation compared to feedforward counterparts, in line with the aforementioned advantages. We provide a general feedback based learning architecture, instantiated using existing RNNs, with the endpoint results on par or better than existing feedforward networks and the addition of the above advantages.

1. Introduction

Feedback is defined to occur when the (full or partial) output of a system is routed back into the input as part of an iterative cause-and-effect process [13]. Utilizing feedback is a strong way of making predictions in various fields, ranging from control theory to psychology [34, 44, 2]. Employing feedback connections is also heavily exercised by the brain [22, 47, 47, 8, 35] suggesting a core role for it in complex cognition. In this paper, we establish that a feedback based learning approach has several core advantages over the commonly employed feedforward paradigm making it a worthwhile alternative. These advantages (elaborated below) are mainly attributed to the fact that the final prediction is made in an iterative, rather than one-time, manner along with an explicit notion of the thus-far output per iteration.

Early Predictions: One advantage is providing early

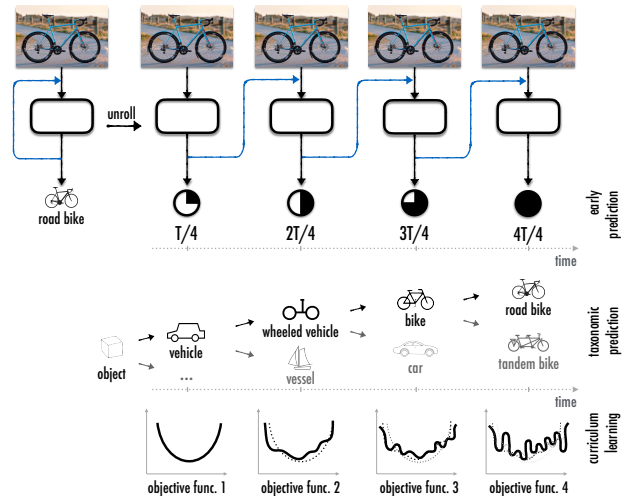*Authors contributed equally.

Figure 1. **A feedback based learning model.** The basic idea is to make predictions in an iterative manner based on a notion of the thus-far outcome. This provides several core advantages: I. enabling early predictions (given total inference time T, early predictions are made in fractions of T); II. naturally conforming to a taxonomy in the output space; and III. better grounds for curriculum learning.

predictions, meaning that estimations of the output can be provided in a fraction of total inference time. This is schematically illustrated in Fig. 1. This property is a result of iterative inference and is in contrast to feedforward where a one-time output is provided only when the signal reaches the end of the network. This is of particular importance in practical scenarios, such as robotics or autonomous driving; e.g. imagine a self driving car that receives a cautionary heads up about possibly approaching a pedestrian on a highway, without needing to wait for the final definite output. Such scenarios are abundant in practice as usually time is crucial and limited computation resources can be reallocated based on early predictions on-the-fly, given a proper uncertainty measure, such as Minimum Bayes Risk [33].

Taxonomy Compliance: Another advantage is making predictions that naturally conform to a hierarchical structure in the output space, e.g. a taxonomy, even when not trained using the taxonomy. The early predictions of the feedback model conform to a coarse classification, while

the later iterations further decompose the coarse class into finer classes. This is illustrated in Fig.1. This is again due to the fact that the predictions happen in an iterative manner coupled with a *coarse-to-fine representation*. The coarse-to-fine representation is naturally developed as the network is forced to make a prediction as early as the first iteration and iteratively improve it in all following iterations.

Episodic Curriculum Learning: The previous advantage is closely related to the concept of Curriculum Learning [4], where gradually increasing the complexity of the task leads to a better training [12, 4, 32]. For non-convex training criteria (such as in ConvNets), a curriculum is known to assist with finding better minima; in convex cases, it improves the convergence speed [4].

As prediction in a feedforward network happens in a one-time manner, a curriculum has to be enforced through feeding the training data in an order based on complexity (i.e. first epochs formed of easy examples and later the hard ones). In contrast, the predictions in a feedback model are made in an iterative form; this enables enforcing a curriculum *through the episodes of prediction for one query*. We call this *Episodic Curriculum Learning*. In other words, sequential easy-to-hard predictions can be enforced for one datapoint (e.g., training the early episodes to predict the species and the later episodes the particular breed). Hence, any taxonomy can be used as a curriculum strategy.

In our model, we define feedback based prediction as a recurrent (weight) shared operation, where at each iteration, the output is estimated and passed onto the next iteration through a hidden state. The next iteration then makes an updated prediction using the shared operation and received hidden state. It is crucial for the hidden state to carry a direction notion of output, otherwise the entire system would be a feedforward pass realized through a recurrent operation [37]. Therefore, we train the network to make a prediction at each iteration by backpropagating the loss in all iteration. We employ a generic architecture for such networks, instantiated using existing RNNs, and empirically prove the aforementioned advantages on various datasets. Though we show that the feedback approach achieves competent final results, the primary goal of this paper is to establish the aforementioned conceptual properties, rather than optimizing for endpoint performance on any benchmark. The developed architectures and pretrained models are available at http://feedbacknet.stanford.edu/.

2. Related Work

There is a notable amount of prior research in machine learning [58, 45, 56, 43, 59, 16, 17, 61, 51, 15, 5, 50] and neuroscience [14, 25, 64] that have commonalities with feedback based learning. We provide a categorized overview of some of the most related works.

Conventional feedforward networks, e.g. AlexNet [31], do not employ recurrence or feedback mechanisms. A number of recent successful methods used recurrence-inspired mechanisms in feedforward models. An example is ResNet [19], introducing parallel residual connections, as well as hypernetworks [18], highway networks [53], stochastic depth [24], RCNN [37], GoogLeNet [55]. These methods are still feedforward srative injection of thus-far output into the system is essential for forming a proper feedback. We empirically show that the feedback mechanism, besides recurrence, is indeed critical (Table 4).

Several recent methods explicitly employed feedback connections [7, 3, 66, 36, 38, 27] with promising results for their task of interest. The majority of these methods are either task specific and/or model temporal problems. Here we put forth and investigate the core advantages of a general feedback based inference. We should also emphasize that feedback in our model is always in the hidden space. This allows us to develop generic feedback based architectures without the requirement of task-specific error-to-input functions [7] (See supplementary material (Sec. 2) for more discussions). Stacked inference methods are also another group of related works [63, 62, 58, 57, 46]. Unlike the method studied here, many of them treat their outputs in isolation and/or do no employ weight sharing.

Another family of methods use feedback like mechanisms for spatial attention [67, 6, 41, 41, 60, 54]. This is usually used for better modeling of long term dependencies, computational efficiency, and spatial localization. Lastly, it is worth noting that Curriculum Learning [12, 32, 4] and making predictions on a taxonomy [23, 52, 9, 11, 28] are well investigated in the literature, though none provided a feedback based approach which is our focus.

3. Feedback Networks

Feedback based prediction has two requirements: (1) iterativeness and (2) rerouting a notion of posterior (output) back into the system in each iteration. We instantiate this by adopting a convolutional recurrent neural network model and connecting the loss to each iteration. The overall process can be summarized as: the image undergoes a shared convolutional operation repeatedly and a prediction is made at each time; the recurrent convolutional operations are trained to produce the best output at each iteration given a hidden state that carries a direction notation of thus-far output. This is depicted in Fig. 2.

3.1. Convolutional LSTM Formulation

In this section, we share the details of our feedback model which is based on stacking a flexible variant of ConvLSTM [66] modules that essentially replace the operations in an LSTM [21] cell with convolutional structures[1]. An LSTM cell uses hidden states to pass information through

---

[1] See supplementary material (Sec. 7) for a discussion on alternatives to LSTM for this purpose, including GRU, vanilla RNN, and ablated LSTM.

Figure 2. **Illustration of our core feedback model and skip connections (shown in red) when unrolled in time.** 'ConvLSTM' and 'L' boxes represent convolutional operations and iteration losses, respectively.

iterations. We briefly describe the connections between stacked ConvLSTMs and the gates in them:

We parametrize the temporal order (i.e. iterations) with time $t = 0, 1, ..., T$ and spatial order of a ConvLSTM module in the stack with depth $d = 0, 1, ..., D$. At depth d and time t, the output of a ConvLSTM module is based on spatial input ($\mathbf{X}_t^{d-1}$), temporal hidden state input ($\mathbf{H}_{t-1}^d$), and temporal cell gate input ($\mathbf{C}_{t-1}^d$).

To compute the output of a ConvLSTM module, the input gate i_t^d and forget gate f_t^d are used to control the information passing between hidden states:

$$i_t^d = \sigma(W_{d,xi}(\mathbf{X}_t^{d-1}) + W_{d,hi}(\mathbf{H}_{t-1}^d)),$$
$$f_t^d = \sigma(W_{d,xf}(\mathbf{X}_t^{d-1}) + W_{d,hf}(\mathbf{H}_{t-1}^d)), \tag{1}$$

where σ is sigmoid function. W is a set of feedforward convolutional operations applied to $\mathbf{X}$ and $\mathbf{H}$. Here W is parametrized by d but not t since the weights of convolutional filters are shared in the temporal dimension. The architecture of W is a design choice and is the primary difference between our ConvLSTM module and Xingjian et al. [66] as we use multilayer convolutional operations for W with flexibility of including residual connections. The depth of W (i.e. the physical depth of a ConvLSTM module) is discussed in Sec. 3.2.

The cell gate $\mathbf{C}_t^d$ is computed as follows:

$$\tilde{C}_t^d = tanh(W_{d,xc}(\mathbf{X}_t^{d-1}) + W_{d,hc}(\mathbf{H}_{t-1}^d)),$$
$$\mathbf{C}_t^d = f_t^d \circ \mathbf{C}_{t-1}^d + i_t^d \circ \tilde{C}_t^d. \tag{2}$$

Finally, the hidden state $\mathbf{H}_t^d$ and output $\mathbf{X}_t^d$ are updated according to the output state o_t and cell state $\mathbf{C}_t^d$:

$$o_t^d = \sigma(W_{d,xo}(\mathbf{X}_t^{d-1}) + W_{d,ho}(\mathbf{H}_{t-1}^d)),$$
$$\mathbf{H}_t^d = o_t^d \circ tanh(\mathbf{C}_t^d), \tag{3}$$
$$\mathbf{X}_t^d = \mathbf{H}_t^d,$$

where '$\circ$' denotes the Hadamard product. Also, we apply batch normalization [26] to each convolutional operation.

For every iteration, loss is connected to the output of the last ConvLSTM module in physical depth. Here, the post processes of ConvLSTM module's output (pooling, fully connected layer, etc.) are ignored for sake of simplicity. $\mathbf{L_t}$ is the cross entropy loss at time t, while C denotes the correct target class number and $\mathbf{L}$ is the overall loss:

$$\mathbf{L} = \sum_{t=1}^{T} \gamma^t \mathbf{L_t}, \text{ where } \mathbf{L_t} = -log\frac{e^{\mathbf{H}_t^D[C]}}{\sum_j e^{\mathbf{H}_t^D[j]}}. \tag{4}$$

γ is a constant discount factor determining the worth of early vs later predictions; we set $\gamma = 1$ in our experiments which gives equal worth to all iterations. [2]

Connecting the loss to all iterations forces the network to attempt the entire task at each iteration and pass the output via the proxy of hidden state (Eq. 4) to future iterations. Thus, the network cannot adopt a representation scheme like feedforward networks that go from low-level (e.g. edges) to high-level representations as merely low-level representations would not be sufficient for accomplishing the whole classification task in early iterations. Instead, the network forms a representation across iterations in a coarse-to-fine manner (further discussed in sections 4.2.2, 4.2.3, and supplementary material's Sec. 3).

We initialize all $\mathbf{X_t^0}$ as the inout image inp, and all $\mathbf{H_0^d}$ as 0, i.e. $\forall t \in \{1, 2, \cdots, T\} : \mathbf{X}_t^0 := inp$ and $\forall d \in \{1, 2, \cdots, D\} : \mathbf{H}_0^d := 0$. The operation of the ConvLSTM module above can be referred to using the simplified notation $\mathfrak{F}(\mathbf{X}_t^{d-1}, \mathbf{H}_{t-1}^d)$.

3.2. Feedback Module Length

We can stack multiple ConvLSTM modules, each a different number of feedforward layers. We categorize feedback networks according to the number of feedforward layers (Conv + BN) within one ConvLSTM module, i.e. the local length of feedback. This is shown in Fig. 3 where the models are named Stack-1, Stack-2, and Stack-All. For Stack-i, i feedforward layers are stacked within one ConvLSTM module. This essentially determines how distributed the propagation of hidden state throughout the network should be (e.g. for the physical depth $\mathcal{D}$, Stack-All architecture would have one hidden state while Stack-1 would have $\mathcal{D}$ hidden states). See supplementary material (Sec. 2) for more discussions. Which length i to pick is a design choice; we provide an empirical study on this in Sec. 4.2.1.

3.3. Temporal Skip Connection

In order to regulate the flow of signal through the network, we include identity skip connections. This was in-

---

[2] **Predicting the 'absolute output' vs an 'adjustment' value**: In this formulation, the absolute output is predicted at each iteration. An alternative would be to predict an 'adjustment value' at each iteration that, when summed with previous iteration's output, would yield the updated absolute output. This approach would have the disadvantage of being applicable to only output spaces with a numerical structure, e.g. regression problems. Problems without a numerical, e.g. classification, or structured space cannot not be solved using this approach.
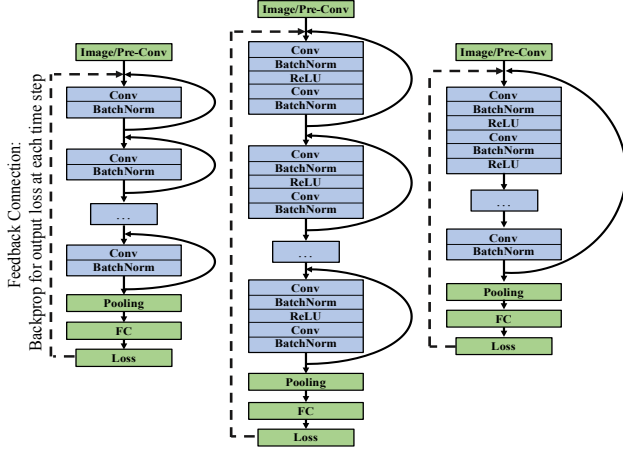
Figure 3. **Feedback networks with different feedback module (ConvL-STM) lengths**. Left, middle, and right show Stack-1, Stack-2, and Stack-All, respectively.

Figure 4. **Computation graph of Feedback vs Feedforward.** $\mathbf{X}_i^j$ denotes the representation at temporal iteration i and physical depth j. Skip connections are not shown for simplicity.

spired by conceptually similar mechanisms, such as the residual connection of ResNet [19] and the recurrent skip coefficients in [69]. The skip connections adopted in the feedback model can be formulated as: with the new input at time t being $\hat{\mathbf{X}}_t^d = \mathbf{X}_t^d + \mathbf{H}_{t-n}^d$, the final representation will be $\mathfrak{F}(\hat{\mathbf{X}}_t^d, \mathbf{H}_{t-n}^d, \mathbf{H}_{t-1}^d)$, where n is the skip length. The skip connections are shown in Fig.2 denoted by the red dashed lines. We set $n = 2$ in our experiments.

Besides regulating the flow, Table 1 quantifies the end-point performance improvement made by such skip connections on CIFAR100 [30] using Stack-2 architecture with physical depth 4 and 8 iterations.

Feedback Net	Top1	Top5
w/o skip connections	67.37	89.97
w/ skip connections	**67.83**	**90.12**

Table 1. **Impact of skip connections in time on CIFAR100** [30]

3.4. Taxonomic Prediction

It is of particular practical value if the predictions of a model conform to a taxonomy. That is, making a correct coarse prediction about a query, if a correct fine prediction cannot be made. Given a taxonomy on the labels (e.g., ImageNet or CIFAR100 taxonomies), we can examine a network's capacity in making taxonomic predictions based on the fine class's Softmax distribution. The probability of a query belonging to the fine class y_i is defined in Softmax as $P(y_i|x; W) = \frac{e^{f_{y_i}}}{\sum_j e^{f_j}}$ for a network with weights W. The probability of a query belonging to the k^{th} higher level coarse class Y_k consisting of $\{y_1, y_2, ..., y_n\}$ is thus the sum of probability of the query being in each of the fine classes:

$$P(Y_k|x; W) = \sum_{i \in 1:n} P(y_i|x; W) = \frac{\sum_{i \in 1:n} e^{f_{y_i}}}{\sum_j e^{f_j}}. \quad (5)$$

Therefore, we use a mapping matrix M, where $M(i, k) = 1$ if $y_i \in Y_k$, to transform fine class distribution to coarse

class distribution. This also gives us the loss for coarse prediction L^{Coarse}, and thus, a coarse prediction p_c is obtained through the fine prediction p_f. In Sec. 4.2.3, it will be shown that the outputs of the feedback network conform to a taxonomy especially in early predictions.

3.5. Episodic Curriculum Learning

As discussed in Sec. 1, the feedback network provides a new way for enforcing a curriculum in learning and enables using a taxonomy as a curriculum strategy. We adopt an iteration-varying loss to enforce the curriculum. We use an annealed loss function at each time step of our k-iteration feedback network, where the relationship of coarse class losses L_t^{Coarse} and fine class losses L_t^{Fine} parametrized by time t is formulated as:

$$L(t) = \zeta L_t^{Coarset} + (1 - \zeta)L_t^{Fine}, \quad (6)$$

where ζ is the weights that balance the contribution of coarse and fine losses. We adopt a linear decay as $\zeta = \frac{t}{k}$, where $t = 0, 1, ..., k$, and k is the end iteration of decaying.

For object classification, the time varying loss function encourages the network to recognize objects in a first coarse then fine manner, i.e. the network learns from the root of an taxonomy tree to its leaves. In Sec. 4.2.4, it will be empirically shown that the feedback based approach well utilizes this curriculum strategy.

3.6. Computation Graph Analysis

Under proper hardware, feedback model also has an advantage on speed over feedforward. This is because a feedback network is a better fit for parallelism compared to feedforward due to having a shallower computation graph (shown in Fig. 4). In the interest of space, we give the full discussion and derivation of the computation graphs in supplementary material (Sec. 4) and only compare their depths here. The computation graph depth of feedforward model with depth $\mathcal{D}$ and that of feedback model with same virtual depth (consisting of m temporal iterations and physical depth n, $\mathcal{D} = m \times n$, and Stack-1 configuration) are $d_{ff} = \mathcal{D} - 1 = mn - 1$ and $d_{fb} = m + n - 1$, respectilvey.

Under a proper hardware scenario where one can do parallel computations to a sufficient extent, inference time can be well measured by the longest distance from root to target (i.e. graph's depth). Therefore, the total prediction time of feedforward network is larger than feedback network's as $d_{ff} = mn - 1 > m + n - 1 = d_{fb}$. Please see supplementary material (Sec. 4) for the depth comparison for early predictions, Stack-i configuration, and traning time.

4. Experimental Results

Our experimental evaluations performed on the three benchmarks of CIFAR100 [30], Stanford Cars [29], and MPII Human Pose [1], are provided in this section.

4.1. Baselines and Terminology

Below we define our terminology and baselines:
Physical Depth: the depth of the convolutional layers from input layer to output layer. For feedback networks, this represents the number of stacked physical layers across all ConvLSTM modules ignoring the temporal dimension.
Virtual Depth: physical depth × number of iterations. This is the effective depth considering both spatial and temporal dimensions. (not applicable to feedforward models.)
Baseline Models: We compare with ResNet[19] and VGG[48] as two of the most commonly used feedforward models and with closest architecture to our convolutional layers. Both baselines have the same architecture, except for the residual connection. We use the same physical module architecture for our method and the baselines. We also compare with ResNet original authors' architecture [19]. The kernel sizes and transitions of filter numbers remain the same as original paper's. In Sec. 4.4, we compare with feedforward Hourglass [42] by making a feedback Hourglass.
Auxiliary prediction layer (aux loss): Feedfoward baselines do not make episodic or mid-network predictions. In order to have a feedforward based baseline for such predictions, we train new pooling→FC→loss layers for different depths of the feedforward baselines (one dedicated aux layers for each desired depth). This allows us to make predictions using the mid-network representations. We train these aux layers by taking the fully trained feedforward network and training the aux layers from shallowest to deepest layer while freezing the convolutional weights.

4.2. CIFAR-100 and Analysis

CIFAR100 includes 100 classes containing 600 images each. The 100 classes (fine level) are categorized into 20 classes (coarse level), forming a 2-level taxonomy. All of the reported quantitative and qualitative results were generated using the fine-only loss (i.e. the typical 100-way classification of CIFAR100), unless specifically mentioned curriculum learning or coarse+fine loss (Eq. 6) were used.

4.2.1 Feedback Module Length

Table 2 provides the results of feedback module length study per the discussion in Sec. 3.2. The physical depth and iteration count are kept constant (physical depth 4 and 4 iterations) for all models. The best performance is achieved when the local feedback length is neither too short nor too long. We found this observation to be valid across different tests and architectures, though the optimal length may not always be 2. In the rest of the experiments for different physical depths, we optimize the value of this hyperparameter empirically (often ends up as 2 or 3). See supplementary material's Sec. 6 for an experimental discussions on the trade-off between physical depth and iteration count as well as optimal iteration number.

Feedback Type	Top1	Top5
Stack-1	66.29	89.58
Stack-2	**67.83**	**90.12**
Stack-All	65.85	89.04

Table 2. **Comparison of different feedback module lengths,** all models have the same physical depth 4 and virtual depth 16.

4.2.2 Early Prediction

We evaluate early predictions of various networks in this section. We conduct this study using a feedback network with virtual depth 32 (similar trends achieved with other depths) and compare it with various feedforward networks. As shown in Fig. 5, at virtual depths of 8, 12, and 16, the feedback network already achieves satisfactory and increasing accuracies. The solid blue and green curves denote the basic feedforward networks with 32 layers; their rightmost performance is their endpoint results, while their early predictions are made using their final pooling→FC→loss layer but applied on mid-network representations. The dashed blue and green curves show the same, with the difference that the trained pooling→FC→loss layers (aux loss, described in Sec. 4.1) are employed for making early predictions. The plot shows that the feedforward networks perform poorly when using their first few layers' representations, confirming that the features learned there are not suitable for completing the ultimate output task (expected) [68]. This is aligned with the hypothesis that feedback model forms its representation in a different and coarse-to-fine manner (further discussed in Sec. 4.2.3).

We also attempted full training and fine tuning the feedforward networks with aux losses, but this never led to a better performance than the reported curves in Fig. 5 by sacrificing either early or endpoint performances. The best results were (comparable to curves in Fig. 5): 6.8%, 10.2%, 13.1%, 13.0%, 59.8%, 66.3%, 68.5% for depths 8, 12, 16, 20, 24, 28, and 32, respectively.
Comparison with Feedforward Ensemble: Although it is memory inefficient and wasteful in training, one can also achieve the effect of early prediction through an ensemble of feedforward models in parallel (i.e. for every depth at which one desires a prediction, have a dedicated feed-

Figure 5. **Evaluation of early predictions.** Comparison of accuracy of feedback (FB) model and feedforward (FF) baselines (ResNet & VGG, with or without auxiliary loss layers)


Figure 6. **Evalaution of taxonomoy based prediction for feedback (FB) and feedforward (FF) networks trained with or without auxiliary layers.** We use only fine loss to train, except for the curriculum learned one.

forward network of that depth). Since running an ensemble of ResNets in parallel has similar optimal hardware requirements of Sec. 3.6, we make a comparison under the same analysis: applying the computation graph depth analysis to a 48 layer virtual depth feedback model (physical depth 12, Stack-3, 4 iterations), if we denote the time to finish one layer of convolution as T, then we have i^{th} iteration result at: $t_i = (12 + 3i)T$. Then the first to last iterations' results (virtual depths $12, 24, 36, 48$) will become available at $12T, 15T, 18T$, and $21T$. To have ResNet results at the same times, we need an ensemble of ResNets with depths $12, 15, 18, 21$. The performance comparison between feedback network and the ensemble is provided in Table 3, showing the advantage of feedback networks.

	Time Steps			
Model	12T	15T	18T	21T
Feedback Network	**67.94**	**70.57**	**71.09**	**71.12**
ResNet Ensemble	66.35	67.52	67.87	68.2

Table 3. **Top1 accuracy comparison between Feedback Net and an ensemble of ResNets** that produce early predictions at the same computation graph depth time steps.

Feedback vs No Feedback: To examine whether the offered observations are caused by feedback or only the recurrence mechanism, we performed a test by disconnecting the loss from all iterations except the last, thus making the model recurrent-feedforward. As shown in Table 4, making the model recurrent-feedforward takes away the ability to make early and taxonomic predictions (discussed next).

	Virtual Depth			
Model	12	24	36	48
Feedback	**67.94**	**70.57**	**71.09**	71.12
Feedback Disconnected (Recurrent Feedforward)	36.23	62.14	67.99	**71.34**

Table 4. **The impact of feedback** on CIFAR100 for a model with virtual depth 48 and four iterations.

Figure 7. **Qualitative results of classification on CIFAR100.** Each row shows a query along with nearest neighbors at different depths for feedback and feedfowrad networks. Orange, blue, and gray represent 'correct fine class', 'correct coarse class but wrong fine class', and 'both incorrect', respectively. Two bottom queries are representative failure cases.

4.2.3 Taxonomic Prediction

We measure the capacity $F(N)$ of network N in making taxonomic predictions (taxonomy compliance) as: the probability of making a correct coarse prediction for a query if it made a wrong fine prediction for it; in other words, how effective it can correct its wrong fine class prediction to a correct coarse class: $F(N) = P(correct(p_c)|!correct(p_f); N)$. As defined in Sec. 3.4, pc and pf stand for coarse and fine prediction, respectively.

The quantitative and qualitative results are provided in Figures 6, 7, and 8. Note that all of these results were naturally achieved, i.e. using fine-only loss and no taxon-
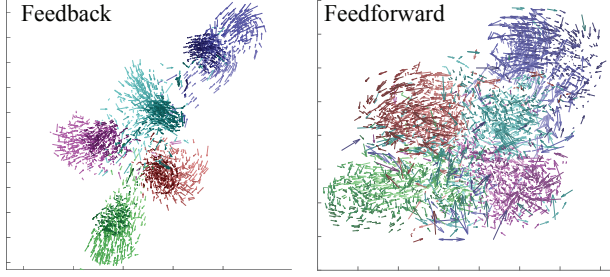
Figure 8. **Timed-tSNE plots showing how the representation evolves through depth/iterations (i.e. how a datapoint moved in representation space)** for each method, on five random classes of CIFAR100. The lighter the hue of the arrow, the earlier the depth/iteration. Feedback's representation is relatively disentangled throughout, while feedforward's representation gets disentangled only towards the end. (Best see on screen. Vector lengths are shown in half to avoid cluttering.)

omy or curriculum learning was used during training (except for the dashed red curve which was trained using curriculum learning; Sec. 4.2.4). Fig. 6 shows feedback network's predictions better complies with a taxonomy even at shallow virtual depths, while feedforward model does not achieve the same performance till the last layer, even when using dedicated auxiliary layers. This is again aligned with the hypothesis that the feedback based approach develops a coarse-to-fine representation and is observed in both figures 7 and 8. In Fig. 7, early prediction classes and nearest neighbor images (using the network representations) for both feedback and feedforward networks are provided, showing significantly more relevant and interpretable early results for feedback.

Timed-tSNE: In Fig. 8, we provide a variant of tSNE [40] plot which we call *timed-tSNE*. It illustrates *how the representation of a network evolves* throughout depth/iterations, when viewed through the window of class labels. For each datapoint, we form a temporally regulated trajectory by connecting a set of 2D tSNE embedding locations. For feedback network, the embeddings of one datapoint come from the representation at different iterations (i.e. i embeddings for a network with i iterations). For feedforward, embeddings come from difference layers. More details provided in supplementary material (Sec. 5).

Fig. 8 suggests that feedforward representation is intertwined at early layers and disentangles the classes only in the last few layers, while feedback's representation is disentangled early on and the updates are mostly around forming fine separation regions. This again supports the hypothesis that feedback develops a coarse-to-fine representation. We also provide activation maps of feedback vs feedforward models in supplementary material (Sec. 5.2) exhibiting notably dissimilar patterns, and thus, dissimilar representations, thought their endpoint numerical results are close.

4.2.4 Curriculum Learning

Table 5 compares the performance of the networks when trained with the fine-only loss vs the episodic coarse-to-

fine curriculum loss (Sec. 3.5). We employed the same episodic curriculum training for the feedback network and the baselines "w/ Aux loss", while the baselines "w/o Aux loss" had to use conventional curriculum training (datapoint sorting) [4]. The best performance with highest boost is achieved by feedback network when using curriculum learning. Also, using the episodic curriculum training improves taxonomic prediction results as shown by the curriculum curve in Fig. 6.

Model	CL	Top1(%)-Fine	Top1(%)-Coarse
Feedback Net	N	68.21	79.7
	Y	**69.57**(+1.34%)	**80.81**(+1.11%)
Feedforward	N	69.36	80.29
ResNet w/ Aux loss	Y	69.24(-0.12%)	80.20(-0.09%)
Feedforward	N	69.36	80.29
ResNet w/o Aux loss	Y	65.69(-3.67%)	76.94(-3.35%)
Feedforward	N	63.56	75.32
VGG w/ Aux loss	Y	64.62(+1.06%)	77.18(+1.86%)
Feedforward	N	63.56	75.32
VGG w/o Aux loss	Y	63.2(-0.36%)	74.97(-0.35%)

Table 5. **Evaluation of the impact of Curriculum Learning (CL)** on CIFAR100. The CL column denotes if curriculum learning was used. The difference made by curriculum for each method is shown in parentheses.

4.2.5 Endpoint Performance Comparison

Table 6 compares the endpoint performance of various feedforward and feedback models on CIFAR100. The detailed architecture of each model is provided in the end of this section. Feedback networks outperform the baselines with the same physical depth by a large margin and work better than or on part with baselines with the same virtual depth or deeper. This ensures that the discussed advantages in early and taxonomic prediction were not achieved at the expense of sacrificing the endpoint performance.

The bottom part of Table 6 shows several recent methods that are not comparable to ours, as they employ additional mechanisms (e.g. stochasticity in depth [24]) which we did not implement in our model. Such mechanisms are independent of feedback and could be used concurrently with it, in the future. However, we include them for the sake of completeness.

Architectures: The detailed architectures of feedback and feedforward networks are: [3]

• Recurrent Block: $Iterate(fi, fo, k, s, n, t)$ denotes our convLSTM recurrent module (defined in Sec. 3.1) which iterates t times and has gate functions, i.e. W, with the feedforward architecture:

$$\to C(fi, fo, k, s) \to BR \to \{C(fo, fo, k, 1) \to BR\}^{n-1}.$$

We denote stacking using $\{...\}^n$ indicating that the module

---

[3]The following naming convention is used: $C(fi, fo, k, s)$: fi input and fo output convolutional filters, kernel size $k \times k$, stride s. $ReLU$: rectified linear unit. BN: batch normalization. $BR = BN + ReLU$. $Avg(k, s)$: average pooling with spatial size $k \times k$, and stride s. $FC(fi, fo)$: fully connected layer with fi inputs, and fo outputs.

Model	Physical Depth	Virtual Depth	Top1 (%)	Top5 (%)
Feedback Net	12	48	**71.12**	**91.51**
(icon)	8	32	69.57	91.01
	4	16	67.83	90.12
Feedforward (ResNet[19]) (icon)	48	-	70.04	90.96
	32	-	69.36	91.07
	12	-	66.35	90.02
	8	-	64.23	88.95
	128*	-	70.92	91.28
	110*	-	72.06	92.12
	64*	-	71.01	91.48
	48*	-	70.56	91.60
	32*	-	69.58	91.55
Feedforward (VGG[48]) (icon)	48	-	55.08	82.1
	32	-	63.56	88.41
	12	-	64.65	89.26
	8	-	63.91	88.90
Highway [53]	19	-	67.76	-
ResNet v2[20]	1001	-	77.29	-
Stochastic Depth [24]	110	-	75.02	-
SwapOut [49]	32 fat	-	77.28	-
RCNN [37]	4 fat	16	68.25	-

Table 6. **Endpoint performance comparison on CIFAR-100.** Baselines denoted with * are the architecture used in the original ResNet paper.

Model	CL	Fine	Coarse
Feedback Net	N	50.33	74.15
	Y	**53.37**(+3.04%)	**80.7**(+6.55%)
Feedforward ResNet-24	N	49.09	72.60
	Y	50.86(+1.77%)	77.25(+4.65%)
Feedforward VGG-24	N	41.04	67.65
	Y	41.87(+0.83%)	70.23(+2.58%)

Table 7. **Evaluations on Stanford Cars dataset**. The CL column denotes if curriculum learning was employed. All methods have (virtual or physical) depth of 24.

feedforward baselines have depth of 24 and feedback network has physical depth 6 and iteration count 4, following the same design in Sections 4.1 & 4.2.5. Full experimental setup is provided in supplementary material (Sec. 8).

4.4. Human Pose Estimation

We evaluated on the regression task of MPII Human Pose estimation [1] benchmark which consists of 40k samples (28k training, 11k testing). Just like we added feedback to feedforward models for CIFAR100 classification and performed comparisons, we applied feedback to the state of the art MPII model Hourglass [42]. We replaced the sequence of ResNet-like convolutional layers in one stack Hourglass with ConvLSTM, which essentially repalced physical depth with virtual depth, and performed backprobapation at each iteration similar to the discussion in Sec. 3.1 (more details about the architecture provided in supplementary material). The performance comparison in Table 8 shows that the feedback model outperforms the deeper feedforward baseline. We provide more results and comparisons with other feedback based methods [7, 3] on this benchmark in supplementary material (Sec. 9).

Method	Physical Depth	Virtual Depth	PCKh
Feedforward-Hourglass	24	-	77.6
Feedback-Hourglass	4	12	**82.3**

Table 8. **Evaluations on MPII Human Pose Dataset.** PCKh is the standard metric measuring body joint localization accuracy [1].

5. Conclusion

We provided a study on feedback based learning, arguing it is a worthwhile alternative to commonly employed feedforward paradigm with several basic advantages: early prediction, taxonomy compliance, and Episodic Curriculum Learning. We also observed that the feedback based approach develops a coarse-to-fine representation that is meaningfully and considerably different from feedforward representations. This study suggests that it would not be far-fetched to find the useful practices of computer vision lying in a feedback based approach in the near future.

in brackets is stacked n times. We use the same architecture as above for all gates and include residual connections in it.
• Preprocess and Postprocess: across all models, we apply the following pre-process: $Input \rightarrow C(3, 16, 3, 1) \rightarrow BR$ and post-process: $\rightarrow Avg(8, 1) \rightarrow FC(64, 100)$
• Feedback Network with physical depth = 8:
$\rightarrow Iterate(16, 32, 3, 2, 2, 4) \rightarrow Iterate(32, 32, 3, 1, 2, 4)$
$\rightarrow Iterate(32, 64, 3, 2, 2, 4) \rightarrow Iterate(64, 64, 3, 1, 2, 4)$
• Feedback Network with physical depth = 12:
$\rightarrow Iterate(16, 16, 3, 1, 3, 4) \rightarrow Iterate(16, 32, 3, 2, 3, 4)$
$\rightarrow Iterate(32, 64, 3, 2, 3, 4) \rightarrow Iterate(64, 64, 3, 1, 3, 4)$
• Baseline Feedforward models with physical depth = $\mathcal{D}$:
$\rightarrow C(16, 32, 3, 2) \rightarrow BR \rightarrow \{C(32, 32, 3, 1) \rightarrow BR\}^{\frac{\mathcal{D}}{2}-1}$
$\rightarrow C(32, 64, 3, 2) \rightarrow BR \rightarrow \{C(64, 64, 3, 1) \rightarrow BR\}^{\frac{\mathcal{D}}{2}-1}$

4.3. Stanford Cars Dataset

To verify the observations made on CIFAR100 on another dataset, we performed the same set of experiments on Stanford Cars dataset [29]. Evaluations of endpoint performance and curriculum learning are provided in table 7. Early prediction and taxonomic prediction curves are provided in supplementary material (Sections 8.1 and 8.2). The experiments show similar trends to CIFAR100's and duplicate the same observations.

All networks were trained from scratch without fine-tuning pretrained ImageNet [10] models [39] or augmenting the dataset with additional images [65]. To suit the relatively smaller amount of training data in this dataset, we use shallower models for both feedforward and feedback:

References

[1] M. Andriluka, L. Pishchulin, P. Gehler, and B. Schiele. 2d human pose estimation: New benchmark and state of the art analysis. In *2014 IEEE Conference on Computer Vision and Pattern Recognition*, pages 3686–3693. IEEE, 2014.

[2] S. J. Ashford and L. L. Cummings. Feedback as an individual resource: Personal strategies of creating information. *Organizational behavior and human performance*, 32(3):370–398, 1983.

[3] V. Belagiannis and A. Zisserman. Recurrent human pose estimation. *arXiv preprint arXiv:1605.02914*, 2016.

[4] Y. Bengio, J. Louradour, R. Collobert, and J. Weston. Curriculum learning. In *Proceedings of the 26th annual international conference on machine learning*. ACM, 2009.

[5] W. Byeon, T. M. Breuel, F. Raue, and M. Liwicki. Scene labeling with lstm recurrent neural networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3547–3555, 2015.

[6] C. Cao, X. Liu, Y. Yang, Y. Yu, J. Wang, Z. Wang, Y. Huang, L. Wang, C. Huang, W. Xu, et al. Look and think twice: Capturing top-down visual attention with feedback convolutional neural networks. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 2956–2964, 2015.

[7] J. Carreira, P. Agrawal, K. Fragkiadaki, and J. Malik. Human pose estimation with iterative error feedback. *arXiv preprint arXiv:1507.06550*, 2015.

[8] R. M. Cichy, D. Pantazis, and A. Oliva. Resolving human object recognition in space and time. *Nature neuroscience*, 17(3):455–462, 2014.

[9] J. Deng, N. Ding, Y. Jia, A. Frome, K. Murphy, S. Bengio, Y. Li, H. Neven, and H. Adam. Large-scale object classification using label relation graphs. In *European Conference on Computer Vision*, pages 48–64. Springer, 2014.

[10] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei. Imagenet: A large-scale hierarchical image database. In *Computer Vision and Pattern Recognition, 2009. CVPR 2009. IEEE Conference on*, pages 248–255. IEEE, 2009.

[11] N. Ding, J. Deng, K. P. Murphy, and H. Neven. Probabilistic label relation graphs with ising models. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 1161–1169, 2015.

[12] J. L. Elman. Learning and development in neural networks: The importance of starting small. *Cognition*, 48(1):71–99, 1993.

[13] F. A. Ford. *Modeling the environment: an introduction to system dynamics models of environmental systems*. Island Press, 1999.

[14] C. D. Gilbert and M. Sigman. Brain states: top-down influences in sensory processing. *Neuron*, 54(5):677–696, 2007.

[15] G. Gkioxari, A. Toshev, and N. Jaitly. Chained predictions using convolutional neural networks. *arXiv preprint arXiv:1605.02346*, 2016.

[16] K. Greff, R. K. Srivastava, and J. Schmidhuber. Highway and residual networks learn unrolled iterative estimation. *arXiv preprint arXiv:1612.07771*, 2016.

[17] K. Gregor and Y. LeCun. Learning fast approximations of sparse coding. In *Proceedings of the 27th International Conference on Machine Learning (ICML-10)*, pages 399–406, 2010.

[18] D. Ha, A. Dai, and Q. V. Le. Hypernetworks. *arXiv preprint arXiv:1609.09106*, 2016.

[19] K. He, X. Zhang, S. Ren, and J. Sun. Deep residual learning for image recognition. *arXiv preprint arXiv:1512.03385*, 2015.

[20] K. He, X. Zhang, S. Ren, and J. Sun. Identity mappings in deep residual networks. *arXiv preprint arXiv:1603.05027*, 2016.

[21] S. Hochreiter and J. Schmidhuber. Long short-term memory. *Neural Computation*, 9(8):1735–1780, Nov. 1997.

[22] C. S. Holling. Resilience and stability of ecological systems. *Annual review of ecology and systematics*, pages 1–23, 1973.

[23] H. Hu, G.-T. Zhou, Z. Deng, Z. Liao, and G. Mori. Learning structured inference neural networks with label relations. *arXiv preprint arXiv:1511.05616*, 2015.

[24] G. Huang, Y. Sun, Z. Liu, D. Sedra, and K. Weinberger. Deep networks with stochastic depth. *arXiv preprint arXiv:1603.09382*, 2016.

[25] J. Hupé, A. James, B. Payne, S. Lomber, P. Girard, and J. Bullier. Cortical feedback improves discrimination between figure and background by v1, v2 and v3 neurons. *Nature*, 394(6695):784–787, 1998.

[26] S. Ioffe and C. Szegedy. Batch normalization: Accelerating deep network training by reducing internal covariate shift. *arXiv preprint arXiv:1502.03167*, 2015.

[27] A. Karpathy and L. Fei-Fei. Deep visual-semantic alignments for generating image descriptions. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3128–3137, 2015.

[28] D. Koller and N. Friedman. *Probabilistic graphical models: principles and techniques*. MIT press, 2009.

[29] J. Krause, M. Stark, J. Deng, and L. Fei-Fei. 3d object representations for fine-grained categorization. In *Proceedings of the IEEE International Conference on Computer Vision Workshops*, pages 554–561, 2013.

[30] A. Krizhevsky and G. Hinton. Learning multiple layers of features from tiny images. 2009.

[31] A. Krizhevsky, I. Sutskever, and G. E. Hinton. Imagenet classification with deep convolutional neural networks. In *Advances in Neural Information Processing Systems*, 2012.

[32] K. A. Krueger and P. Dayan. Flexible shaping: How learning in small steps helps. *Cognition*, 110(3):380–394, 2009.

[33] S. Kumar and W. Byrne. Minimum bayes-risk decoding for statistical machine translation. Technical report, DTIC Document, 2004.

[34] E. B. Lee and L. Markus. Foundations of optimal control theory. Technical report, DTIC Document, 1967.

[35] T. S. Lee and D. Mumford. Hierarchical bayesian inference in the visual cortex. *JOSA A*, 20(7):1434–1448, 2003.

[36] K. Li, B. Hariharan, and J. Malik. Iterative instance segmentation. *arXiv preprint arXiv:1511.08498*, 2015.

[37] M. Liang and X. Hu. Recurrent convolutional neural network for object recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3367–3375, 2015.

[38] Q. Liao and T. Poggio. Bridging the gaps between residual learning, recurrent neural networks and visual cortex. *arXiv preprint arXiv:1604.03640*, 2016.

[39] T.-Y. Lin, A. RoyChowdhury, and S. Maji. Bilinear cnn models for fine-grained visual recognition. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 1449–1457, 2015.

[40] L. v. d. Maaten and G. Hinton. Visualizing data using t-sne. *Journal of Machine Learning Research*, 9(Nov):2579–2605, 2008.

[41] V. Mnih, N. Heess, A. Graves, et al. Recurrent models of visual attention. In *Advances in neural information processing systems*, pages 2204–2212, 2014.

[42] A. Newell, K. Yang, and J. Deng. Stacked hourglass networks for human pose estimation. *arXiv preprint arXiv:1603.06937*, 2016.

[43] M. Oberweger, P. Wohlhart, and V. Lepetit. Training a feedback loop for hand pose estimation. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 3316–3324, 2015.

[44] A. G. Parlos, K. T. Chong, and A. F. Atiya. Application of the recurrent multilayer perceptron in modeling complex process dynamics. *IEEE Transactions on Neural Networks*, 5(2):255–266, 1994.

[45] P. H. Pinheiro and R. Collobert. Recurrent convolutional neural networks for scene labeling. In *ICML*, pages 82–90, 2014.

[46] V. Ramakrishna, D. Munoz, M. Hebert, J. A. Bagnell, and Y. Sheikh. Pose machines: Articulated pose estimation via inference machines. In *European Conference on Computer Vision*, pages 33–47. Springer, 2014.

[47] N. C. Rust and J. J. DiCarlo. Selectivity and tolerance (invariance) both increase as visual information propagates from cortical area v4 to it. *The Journal of Neuroscience*, 30(39):12978–12995, 2010.

[48] K. Simonyan and A. Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014.

[49] S. Singh, D. Hoiem, and D. Forsyth. Swapout: Learning an ensemble of deep architectures. *arXiv preprint arXiv:1605.06465*, 2016.

[50] R. Socher, B. Huval, B. P. Bath, C. D. Manning, and A. Y. Ng. Convolutional-recursive deep learning for 3d object classification. In *NIPS*, volume 3, page 8, 2012.

[51] R. Socher, C. C. Lin, C. Manning, and A. Y. Ng. Parsing natural scenes and natural language with recursive neural networks. In *Proceedings of the 28th international conference on machine learning (ICML-11)*, pages 129–136, 2011.

[52] Y. Song, M. Zhao, J. Yagnik, and X. Wu. Taxonomic classification for web-based videos. In *Computer Vision and Pattern Recognition (CVPR), 2010 IEEE Conference on*, pages 871–878. IEEE, 2010.

[53] R. K. Srivastava, K. Greff, and J. Schmidhuber. Highway networks. *arXiv preprint arXiv:1505.00387*, 2015.

[54] M. F. Stollenga, J. Masci, F. Gomez, and J. Schmidhuber. Deep networks with internal selective attention through feedback connections. In *Advances in Neural Information Processing Systems*, pages 3545–3553, 2014.

[55] C. Szegedy, W. Liu, Y. Jia, P. Sermanet, S. Reed, D. Anguelov, D. Erhan, V. Vanhoucke, and A. Rabinovich. Going deeper with convolutions. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1–9, 2015.

[56] J. Tompson, R. Goroshin, A. Jain, Y. LeCun, and C. Bregler. Efficient object localization using convolutional networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 648–656, 2015.

[57] A. Toshev and C. Szegedy. Deeppose: Human pose estimation via deep neural networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1653–1660, 2014.

[58] Z. Tu. Auto-context and its application to high-level vision tasks. In *Computer Vision and Pattern Recognition, 2008. CVPR 2008. IEEE Conference on*, pages 1–8. IEEE, 2008.

[59] A. Veit, M. J. Wilber, and S. Belongie. Residual networks behave like ensembles of relatively shallow networks. In *Advances in Neural Information Processing Systems*, pages 550–558, 2016.

[60] Q. Wang, J. Zhang, S. Song, and Z. Zhang. Attentional neural network: Feature selection using cognitive feedback. In *Advances in Neural Information Processing Systems*, pages 2033–2041, 2014.

[61] Z. Wang, S. Chang, J. Zhou, M. Wang, and T. S. Huang. Learning a task-specific deep architecture for clustering. In *Proceedings of the 2016 SIAM International Conference on Data Mining*, pages 369–377. SIAM, 2016.

[62] D. J. Weiss and B. Taskar. Structured prediction cascades. In *AISTATS*, pages 916–923, 2010.

[63] D. H. Wolpert. Stacked generalization. *Neural networks*, 5(2):241–259, 1992.

[64] D. Wyatte, T. Curran, and R. O'Reilly. The limits of feedforward vision: Recurrent processing promotes robust object recognition when objects are degraded. *Journal of Cognitive Neuroscience*, 24(11):2248–2261, 2012.

[65] S. Xie, T. Yang, X. Wang, and Y. Lin. Hyper-class augmented and regularized deep learning for fine-grained image classification. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2645–2654, 2015.

[66] S. Xingjian, Z. Chen, H. Wang, D.-Y. Yeung, W.-k. Wong, and W.-c. Woo. Convolutional lstm network: A machine learning approach for precipitation nowcasting. In *Advances in Neural Information Processing Systems*, 2015.

[67] K. Xu, J. Ba, R. Kiros, K. Cho, A. Courville, R. Salakhutdinov, R. S. Zemel, and Y. Bengio. Show, attend and tell: Neural image caption generation with visual attention. *arXiv preprint arXiv:1502.03044*, 2(3):5, 2015.

[68] M. D. Zeiler and R. Fergus. Visualizing and understanding convolutional networks. In *European Conference on Computer Vision*, pages 818–833. Springer, 2014.

[69] S. Zhang, Y. Wu, T. Che, Z. Lin, R. Memisevic, R. Salakhutdinov, and Y. Bengio. Architectural complexity measures of recurrent neural networks. *arXiv preprint arXiv:1602.08210*, 2016.