

Spatial Knowledge Distillation to aid Visual Reasoning

Somak Aditya^{†*}
 Adobe Research (BEL)
 saditya@adobe.com

Rudra Saha*, Yezhou Yang, Chitta Baral
 CIDSE, Arizona State University
 {rsaha3, yz.yang, chitta}@asu.edu

Abstract

For tasks involving language and vision, the current state-of-the-art methods tend not to leverage any additional information that might be present to gather relevant (commonsense) knowledge. A representative task is Visual Question Answering where large diagnostic datasets have been proposed to test a system's capability of answering questions about images. The training data is often accompanied by annotations of individual object properties and spatial locations. In this work, we take a step towards integrating this additional privileged information in the form of spatial knowledge to aid in visual reasoning. We propose a framework that combines recent advances in knowledge distillation (teacher-student framework), relational reasoning and probabilistic logical languages to incorporate such knowledge in existing neural networks for the task of Visual Question Answering. Specifically, for a question posed against an image, we use a probabilistic logical language to encode the spatial knowledge and the spatial understanding about the question in the form of a mask that is directly provided to the teacher network. The student network learns from the ground-truth information as well as the teachers prediction via distillation. We also demonstrate the impact of predicting such a mask inside the teachers network using attention. Empirically, we show that both the methods improve the test accuracy over a state-of-the-art approach on a publicly available dataset.

1. Introduction

Vision and language tasks such as Visual Question Answering (VQA) are often considered as “AI-complete” tasks [1] since they require multi-modal knowledge beyond a single sub-domain. Recently, the VQA1.0 dataset was proposed as a representative dataset for the task of VQA [1]. This task aims to combine efforts from three broad sub-fields of AI namely image understanding, language understanding and reasoning. Despite its popularity, most of its questions focus on object recognition in images and natural

language understanding. Question-Image pairs where a system may require compositional reasoning or reasoning with external knowledge, seem to be largely absent. To explicitly assess the reasoning capability, several specialized datasets have been proposed, that emphasize specifically on questions requiring complex multiple-step reasoning (CLEVR [12], Sort-of-Clevr [24]) or questions that require reasoning using external knowledge (F-VQA [27]).

Current state-of-the-art neural architectures do not explicitly model such external knowledge and reasoning with them to solve visual reasoning tasks. Several researchers [15, 16] in their works have pointed out the necessity of explicit modeling of such knowledge. This necessitates considering the following issues:

- i) *What kind of knowledge is needed?*
- ii) *Where and how to get them?* and
- iii) *What kind of reasoning mechanism to adopt for such knowledge?*

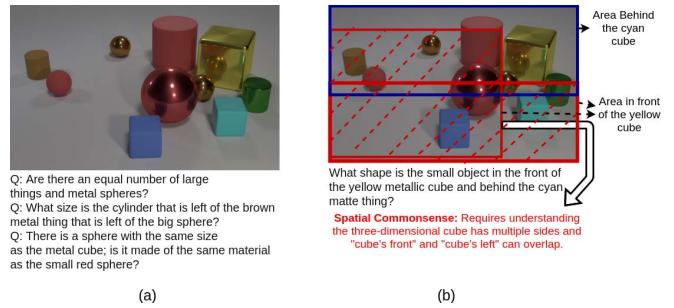


Figure 1: (a) An image and a set of questions from the CLEVR dataset. Questions often require multiple-step reasoning, for example *in the second question, one needs to identify the big sphere, then recognize the reference to the brown metal cube, which then refers to the root object, that is, the brown cylinder.* (b) *An example of spatial commonsense knowledge needed to solve a CLEVR-type question.*

To understand the kind of external knowledge required, we investigate the CLEVR dataset proposed in [12]. This dataset explicitly asks questions that require relational and multi-step reasoning. An example is provided in Fig. 1(a). In this dataset, the authors create synthetic images con-

^{0*} denotes equal contribution. [†]: work was done when author was in CIDSE, ASU.

sisting of a set of objects that are placed randomly within the image. Each object is created randomly by varying its shape, color, size and texture. For each image, 10 complex questions are generated. Each question inquires about an object or a set of objects in the image. To understand which object(s) the question is referring to, one needs to decipher the clues that are provided about the property of the object or the spatial relationships with other objects. This can be a multiple-step process, that is: first recognize object A, that refers to object B, which refers to C and so on. There have been multiple architectures proposed to answer such complex questions. Authors in [9] attempt to learn a structured program from the natural language question. This program acts as a structured query over the objects and relationship information provided as a scene graph and can retrieve the desired answer. More interestingly, the authors in [24] model relational reasoning explicitly in the neural network architecture and propose a generic relational reasoning module to answer questions. This is one of the first known attempt to formulate a differentiable function to embody a generic relational reasoning module that is traditionally formulated using logical reasoning languages. The failure cases depicted by this work, often points to the lack of complex commonsense knowledge such as, *the front of cube should consist of front of all visible side of cubes*. These examples point that spatial commonsense knowledge might help answer questions such as in Fig. 1(b). Even though procuring such knowledge explicitly is difficult, we observe that parsing the questions and additional scene-graph information can help “disambiguate” the area of the image on which a phrase of a question focuses on.

In this work, to the best of our knowledge, we make the first attempt of incorporating such additional information in the form of spatial knowledge into existing architectures for vision and language tasks. Specifically, we concentrate on the task of VQA for questions that require multiple-step (relational) reasoning, and we explore how a recently proposed relational reasoning based architecture [24] can be improved further with the aid of additional spatial knowledge extracted from the image using the question and the scene-graph information. This is an important avenue, as humans often use a large amount of external knowledge to solve tasks that they have acquired through years of experience¹. To extract such spatial knowledge in a form that can be integrated and reasoned with, we take inspiration from techniques from the field of Knowledge Representation and Reasoning (KR&R) and utilize a reasoning engine viz. Probabilistic Soft Logic (PSL) [2]. In practice, these languages and their available implementa-

¹The authors in [15] quoted a reviewer’s comment: “Human learners - unlike DQN and many other deep Learning systems - approach new problems armed with extensive prior experience.”. The authors also ask “How do we bring to bear rich prior knowledge to learn new tasks and solve new problems?”.

tions are often susceptible to the high amount of noises in real-world datasets and hence, their direct applications have been somewhat limited. One can assume that in order to provide robust, interpretable and accurate solutions, one needs to leverage both the robustness and interpretability of declarative logical reasoning languages and the high-level representation learning capability of deep learning.

We rely on the knowledge distillation paradigm [8] in order to integrate the extracted spatial knowledge with the relational reasoning architecture. Knowledge distillation aims to transfer the predictions learned by a complex model, often regarded as a teacher model, to that of a simpler model, usually deemed as a student model, via distillation [8]. There are various flavours of the knowledge distillation paradigm depending upon the complexity of the teacher’s model as well as the amount of information it has access to relative to the student’s model.

Thus to solve our task, we propose a student-teacher network based architecture, where the teacher has access to privileged information. For the VQA, this privileged information is the spatial knowledge required to answer the question in the form of an attentive image mask based on the question and the scene-graph information. The student network is the existing architecture we want to distill this knowledge into. We provide two methods for calculating the mask; i) when the object and relationships are provided for an image, one can calculate a mask using probabilistic reasoning, and ii) if such data is not available, such a mask can be calculated inside the network using attention. We experiment on the Sort-of-Clevr dataset, and empirically show that both these methods outperform a state-of-the-art relational reasoning architecture. We observe that the teacher model (using the spatial knowledge inferred by PSL inference) achieves a sharp 13.7% jump in test accuracy over the baseline architecture. The existing architecture *i.e.* the student model, distilled with this knowledge, shows a generalization boost of 6.2% as well. We also provide ablation studies of the reasoning mechanism on (questions and scene information from) the CLEVR dataset.

2. Related Work

Our work is influenced by the following thrusts of work: probabilistic logical reasoning, spatial reasoning, reasoning in neural networks, knowledge distillation; and the target application area of Visual Question Answering.

Recently, researchers from the KR&R community, and the Probabilistic Reasoning community have come up with several robust probabilistic reasoning languages which are deemed more suitable to reason with real-world noisy data, and incomplete or noisy background knowledge. Some of the popular ones among these reasoning languages are Markov Logic Network [23], Probabilistic Soft Logic [2], and ProbLog [6]. Even though these new theories are con-

siderable large steps towards modeling uncertainty (beyond previous languages engines such as Answer Set Programming [4]); the benefit of using these reasoning engines has not been successfully shown on large real-world datasets. This is one of the reasons, recent advances in deep learning, especially the works of modeling knowledge distillation [8, 25] and relational reasoning have received significant interest from the community.

Modeling of Spatial Knowledge and reasoning using such knowledge in 2D or 3D space has given rise to multiple interesting works in both Computer Vision and Robotics, collectively termed as Qualitative Spatial Reasoning (QSR). Randell *et al.* [22] proposed an interval logic for reasoning about space. Cohn and Renz [5] proposed advancements over previous languages aimed at robotic navigation in 2D or 3D space. In these languages, the relations between two objects are modeled spatially. Our work is also influenced by this series of works (such as Region Connection Calculus etc.), in the sense of what “privileged information” we expect along-with the image and the question. For the CLEVR dataset, the relations `left`, `right`, `front`, `behind` can be used as a closed set of spatial relations among the objects and that often suffices to answer most questions. For real images, a scene graph that encodes spatial relations among objects and regions, such as proposed in [7] would be useful to integrate our methods.

Popular probabilistic reasoning mechanisms from the statistical community often define distribution with respect to Probabilistic Graphical Models. There have been a few attempts to model such graphical models in conjunction with deep learning architectures [31]. However, multi-step relational reasoning, and reasoning with external domain or commonsense knowledge² require the robust structured modeling of the world as adopted by KR&R languages. In its popular form, these reasoning languages often use predicates to describe the current world, such as `color(hair, red)`, `shape(object1, sphere)`, `material(object1, metal)` etc; and then declare rules that the world should satisfy. Using these rules, truth values of unknown predicates are obtained, such as `ans(?x, O)` etc. Similarly, the work in [24], defines the relational reasoning module as $RN(O) = f_{\phi}\left(\sum_{i,j} g_{\theta}(o_i, o_j)\right)$, where O denote all objects. In this work, the relation between a pair of objects (i.e. g_{θ}) and the final function over this collection of relationships i.e. f_{ϕ} are defined as multilayer perceptrons (MLP) and are learnt using gradient descent in an end-to-end manner. This model’s simplicity and its close resemblance to traditional reasoning mechanisms motivates us to pursue further and integrate external knowledge.

²An example of multi-step reasoning: if event A happens, then B will happen. The event B causes action C only if event D does not happen. For reasoning with knowledge: consider for a image with a giraffe, we need to answer “Is the species of the animal in the image and an elephant same?”

Several methods have been proposed to distill knowledge from a larger model to a smaller model or from a model with access to privileged information to a model without such information. Hinton *et al.* [8] first proposed a framework where a large cumbersome model is trained separately and a smaller student network learns from both groundtruth labels and the large network. Independently, Vapnik *et al.* [25] proposed an architecture where the larger (or the teacher) model has access to privileged information and the student model does not. These models together motivated many natural language processing researchers to formulate textual classification tasks as a teacher-student model, where the teacher has privileged information, such as a set of rules; and the student learns from the teacher and the ground-truth data. The imitation parameter controls how much the student *trusts* the teacher’s decision. In [11], an iterative knowledge distillation is proposed where the teacher and the student learn iteratively and the convolutional network’s parameters are shared between the models. In [10], the authors propose to solve sentiment classification, by encoding explicit logical rules and integrating the grounded rules with the teacher network. These applications of teacher-student network only exhibited success with classification problems with very small number of classes (less than three).

In this paper, we show a knowledge distillation integration with privileged information which is applied to a 28-class classification, and we observe that it improves by a large margin on the baseline. In [30], the authors use encoded linguistic knowledge in the form of $P(pred|obj, subj)$ to perform Visual Relationship Detection. In this work, we apply knowledge distillation in a visual question answering setting, that require both visual reasoning and question understanding.

In the absence of the scene information or in cases where such information is expensive to obtain, an attention mask over the image can be predicted inside the network based upon the posed question. Attention mechanism has been successfully applied in image captioning [28, 19], machine translation[3, 26] and visual question answering [29]. In [29], a stacked attention network was used to predict a mask over the image. They use the question vector separately to query specific image features to create the first level of attention. In contrast, we combine the question vector with the whole image features to predict a coarse attention mask.

3. Additional Knowledge Integration Method

In this section, we explain the various components of our proposed framework for integrating additional spatial information with existing neural architectures. We start by formalizing the probabilistic reasoning mechanism which enables us to extract such spatial knowledge in the presence of scene information. Then, we describe the knowledge distillation paradigm [8] that enables us to infuse this ex-

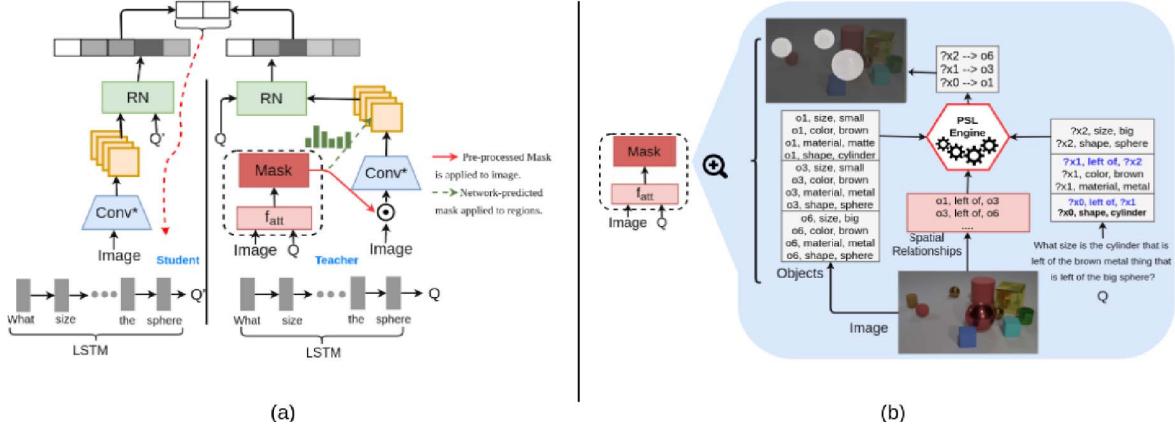


Figure 2: (a) The Teacher-Student Distillation Architecture. As the base of both teacher and student, we use the architecture proposed by the authors in [24]. For the experiment with pre-processed mask generation, we pass a masked image through the convolutional network and for the network-predicted mask, we use the image and question to predict an attention mask over the regions. (b) We show the internal process of mask creation.

tracted knowledge into existing networks which in our case is a relational reasoning architecture [24]. We also outline the in-network computation required in the absence of the scene-graph information.

3.1. Probabilistic Reasoning Mechanism

In order to reason about the spatial relations among the objects in a scene and textual mentions of those objects in the question, we choose Probabilistic Soft Logic (PSL) [2] as our reasoning engine. Using PSL provides us three advantages: i) (Robust Joint Modeling) from the statistical side, PSL models the joint distribution of the random variables using a Hinge-Loss Markov Random Field, ii) (interpretability) we can use clear readable declarative rules that (directly) relates to defining the clique potentials, and iii) (Convex Optimization) the optimization function of PSL is designed in a way so that the underlying function remains convex and that provides an added advantage of faster inference. We use PSL, as it has been successfully used in Vision applications [17] in the past and it is also known to scale up better than its counterparts [23].

3.1.1 Hinge-Loss Markov Random Field and PSL

Hinge-Loss Markov Random Fields (HL-MRF) is a general class of continuous-valued probabilistic graphical model. An HL-MRF is defined as follows: Let \mathbf{y} and \mathbf{x} be two vectors of n and n' random variables respectively, over the domain $D = [0, 1]^{n+n'}$. The feasible set \tilde{D} is a subset of D , which satisfies a set of inequality constraints over the random variables.

A Hinge-Loss Markov Random Field \mathbb{P} is a probability density over D , defined as: if $(\mathbf{y}, \mathbf{x}) \notin \tilde{D}$, then $\mathbb{P}(\mathbf{y}|\mathbf{x}) = 0$; if $(\mathbf{y}, \mathbf{x}) \in \tilde{D}$, then:

$$\mathbb{P}(\mathbf{y}|\mathbf{x}) \propto \exp(-f_{\mathbf{w}}(\mathbf{y}, \mathbf{x})). \quad (1)$$

PSL combines the declarative aspect of reasoning languages with conditional dependency modeling power of undirected graphical models. In PSL a set of weighted if-then rules over first-order predicates is used to specify a Hinge-Loss Markov Random field.

In general, let $\mathcal{C} = (C_1, \dots, C_m)$ be such a collection of weighted rules where each C_j is a disjunction of literals, where each literal is a variable y_i or its negation $\neg y_i$, where $y_i \in \mathbf{y}$. Let I_j^+ (resp. I_j^-) be the set of indices of the variables that are not negated (resp. negated) in C_j . Each C_j can be represented as:

$$w_j : \vee_{i \in I_j^+} y_i \leftarrow \wedge_{i \in I_j^-} y_i, \quad (2)$$

or equivalently, $w_j : \vee_{i \in I_j^-} (\neg y_i) \vee \vee_{i \in I_j^+} y_i$. A rule C_j is associated with a non-negative weight w_j . PSL relaxes the boolean truth values of each ground atom a (constant term or predicate with all variables replaced by constants) to the interval $[0, 1]$, denoted as $V(a)$. To compute soft truth values, Lukasiewicz's relaxation [14] of conjunctions (\wedge), disjunctions (\vee) and negations (\neg) are used:

$$V(l_1 \wedge l_2) = \max\{0, V(l_1) + V(l_2) - 1\}$$

$$V(l_1 \vee l_2) = \min\{1, V(l_1) + V(l_2)\}$$

$$V(\neg l_1) = 1 - V(l_1).$$

In PSL, the ground atoms are considered as random variables, and the joint distribution is modeled using Hinge-Loss Markov Random Field (HL-MRF).

In PSL, the hinge-loss energy function $f_{\mathbf{w}}$ is defined as:

$$f_{\mathbf{w}}(\mathbf{y}) = \sum_{C_j \in \mathcal{C}} w_j \max\left\{1 - \sum_{i \in I_j^+} V(y_i) - \sum_{i \in I_j^-} (1 - V(y_i)), 0\right\}. \quad (3)$$

The maximum-a posteriori (MAP) inference objective of

PSL becomes:

$$\begin{aligned}
\arg \max_{\mathbf{y} \in [0,1]^n} P(\mathbf{y}) &\equiv \arg \max_{\mathbf{y} \in [0,1]^n} \exp(-f_{\mathbf{w}}(\mathbf{y})) \\
&\equiv \arg \min_{\mathbf{y} \in [0,1]^n} \sum_{C_j \in \mathcal{C}} w_j \max \left\{ 1 - \sum_{i \in I_j^+} V(y_i) \right. \\
&\quad \left. - \sum_{i \in I_j^-} (1 - V(y_i)), 0 \right\}, \tag{4}
\end{aligned}$$

where the term $w_j \times \max \left\{ 1 - \sum_{i \in I_j^+} V(y_i) - \sum_{i \in I_j^-} (1 - V(y_i)), 0 \right\}$ measures the “distance to satisfaction” for each grounded rule C_j .

3.2. Knowledge Distillation Framework

While PSL provides a probabilistic knowledge representation, as shown in Fig. 2(b), a mechanism is needed to utilize them under the deep neural networks based systems. We use the generalized knowledge distillation paradigm [18], where the teacher’s network can be a larger network performing additional computation or have access to privileged information, to achieve this integration resulting in two different architectures i) (External Mask) teacher with provided ground-truth mask, ii) (In-Network Mask) teacher predicts the mask with additional computation. Here, we provide general formulations for both methods and give an overview of how the external mask is calculated³.

3.2.1 General Architecture

The general architecture for the teacher-student network is provided in Fig. 2(a). Let us denote the teacher network as q_{ϕ} and the student network as p_{θ} . In both scenarios, the student network uses the relational reasoning network [24] to predict the answer. The teacher network uses an LSTM to process the question, and a convolutional neural network to process the image. Features from the convolutional network and the final output from the LSTM is used as input to the relational reasoning module to predict an answer. Additionally in the teacher network, we predict a mask. For the External Mask setting, the mask is predicted by a reasoning engine and applied to the image, and for the attention setting, the mask is predicted using the image and text features and applied over the output from the convolution. The teacher network q_{ϕ} is trained using softmax cross-entropy loss against the ground truth answers for each question. The student network is trained using knowledge distillation with the following objective:

$$\begin{aligned}
\theta = \arg \min_{\theta \in \Theta} \sum_{n=1}^N &(1 - \pi) \ell_1(\mathbf{y}_n, \sigma_{\theta}(\mathbf{x}_n)) \tag{5} \\
&+ \pi \ell_2(\mathbf{s}_n, \sigma_{\theta}(\mathbf{x}_n)),
\end{aligned}$$

³A detailed example of how we estimate these predicates is provided in Supplementary material.

where \mathbf{x}_n is the image-question pair, and \mathbf{y}_n is the answer that is available during the training phase; the $\sigma_{\theta}(\cdot)$ is the usual softmax function; \mathbf{s}_n is the soft prediction vector of q_{ϕ} on \mathbf{x}_n and ℓ_i denotes the loss functions selected according to specific experiments (usually ℓ_1 is cross-entropy and ℓ_2 is euclidean norm). π is often called the imitation parameter and determines how much the student trusts the teacher’s predictions.

3.2.2 External Mask Prediction

This experimental setting is motivated by the widely available scene graph information in large datasets starting from Sort-of-Clevr and CLEVR to Visual Genome. We use the following information about the objects and their relationships in the image: i) the list of *attribute, value* pairs for each object, ii) the spatial relationships between objects, and iii) each object’s relative location in the image.

We view the problem as a special case of the bipartite matching problem, where there is one set of textual mentions (M) of the actual objects and a second set of actual objects (O). Using probabilistic reasoning we find a matching between object-mention pairs based on how the attribute-value pairs match between the objects and the corresponding mentions, and when mention-pairs are consistently related (such as *larger than*, *left to*, *next to*) as their matched object-pairs. Using the scene graph data, and by parsing the natural language question, we estimate the value of the following predicates: $attr_o(O, A, V)$, $attr_m(M, A, V)$ and $consistent(A, O, O_1, M, M_1)$. The predicate $attr_m(M, A, V)$ denotes the confidence that the value of the attribute A of the textual mention M is V . The predicate $attr_o(O, A, V)$ is similar and denotes a similar confidence for the object O . The predicate $consistent(R, O, O_1, M, M_1)$ indicates the confidence that the textual mentions M and M_1 are consistent based on a relationship R (spatial or attribute based), if M is identified with the object O and M_1 is identified with the object O_1 . Using only these two predicate values, we use the following two rules to estimate which objects relate to which textual mentions.

$$\begin{aligned}
w_1 : candidate(M, O) &\leftarrow object(O) \wedge mention(M) \\
&\quad \wedge attr_o(O, A, V) \wedge attr_m(M, A, V). \\
w_2 : candidate(M, O) &\leftarrow object(O) \wedge mention(M) \\
&\quad \wedge candidate(M, O) \\
&\quad \wedge candidate(M_1, O_1) \\
&\quad \wedge consistent(A, O, O_1, M, M_1).
\end{aligned}$$

We use the grounded rules (variables replaced by constants) to define the clique potentials and use eq. 4 to find the confidence scores of grounded $candidate(M, O)$ predicates. Using this mention to object mapping, we use the objects that the question refers to. For each object, we use the center location, and create a heatmap that decays with distance

from the center. We use a union of these heatmaps as the mask. This results into a set of spherical masks over the objects mentioned in the question, as shown in Fig. 2(b). To validate our calculated masks, we annotate the CLEVR validation set with the ground-truth objects, using the ground-truth structured program. We observe that our PSL-based method can achieve a 75% recall and 70% precision in predicting the ground-truth objects for a question.

In Figure 3, we provide more details of the calculated PSL predicates for the example question and image in Figure 2(b). We use this top collection of objects and their relative locations to create small spherical masks over the relevant objects in the images.

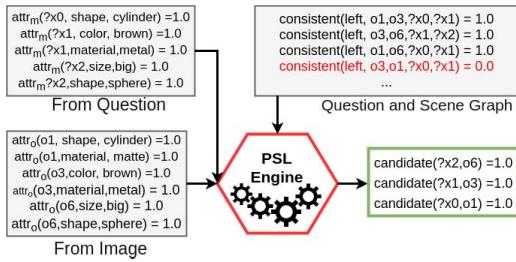


Figure 3: We elaborate on the calculated PSL predicates for the example image and question in Figure 2(b). The underlying optimization benefits from the negative examples (the *consistent* predicate with 0.0, marked in red). Hence, these predicates are also included in the program.

3.2.3 In-Network Mask Prediction

The External Mask setting requires privileged information such as scene graph data about the image, which includes the spatial relations between objects. Such information is often expensive to obtain. Hence, in one of our experiments, we attempt to emulate the mask creation inside the network. We formulate the problem as attention mask generation over image regions using the image ($\mathbf{x}_I \in \mathbb{R}^{64 \times 64 \times 3}$) and the question ($\mathbf{x}_q \in \mathbb{R}^{w \times d}$). The calculation can be summarized by the following equations:

$$\begin{aligned} r_I &= conv^*(\mathbf{x}_I), \quad q_{emb} = LSTM(\mathbf{x}_q). \\ v &= \tanh(W_I r_I + W_q q_{emb} + b). \\ \alpha &= \exp(v) / \sum_{r=1}^{x \times y} \exp(v_r), \end{aligned} \quad (6)$$

where r_I is $x \times y$ regions with o_c output channels, $q_{emb} \in \mathbb{R}^h$ is the final hidden state output from $LSTM$ (hidden state size is h); $W_I (\in \mathbb{R}^{xyo_c \times xy})$ and $W_q (\in \mathbb{R}^{xy \times h})$ are the weights and b is the corresponding bias. Finally, the attention α over regions is obtained by exponentiating the weights and then normalizing them. The attention α is then reshaped and element-wise multiplied with the region features extracted from the image. This is considered as a mask

over the image regions conditioned on the question vector and the image features.

4. Experiments and Results

We propose two architectures, one where the teacher has privileged information and the other where the teacher performs additional calculation using auxiliary in-network modules. We perform experiments to validate whether the direct addition of information (external mask), or additional modules (model with attention) improves the teacher's performance over the baseline. We also perform similar experiments to validate whether this learned knowledge can be distilled to existing neural networks (student model). Additionally, we conduct ablation studies on the probabilistic logical mechanism using which we predict a ground-truth mask from the question and the scene information.

4.1. Setup

As our testbed, we use the “Sort-of-Clevr” from [24] and the CLEVR dataset from [12]. As the original Sort-of-Clevr dataset is not publicly available, we create the synthetic dataset as described by the authors⁴. We use similar specification, i.e., there are 6 objects per image, where each object is either a circle or a rectangle, and we use 6 colors to identify each different object. Unlike the original dataset, we generate natural language questions along with their one-hot vector representation. In our experiments we primarily use the natural-language question. We only use the one-hot vector to replicate results of the baseline Relational Network (RN)⁵. For our experiments, we use 9800 images for training, 200 images each for validation and testing. There are 10 question-answer pairs for each image. For Sort-of-Clevr, we use four convolutional layers with 32, 64, 128 and 256 kernels, ReLU non-linearities, and batch normalization. The questions were passed through an LSTM where the word embeddings are initialized with 50-dimensional Glove embeddings [20]. The LSTM output and the convolutional features are passed through the RN network⁶. The baseline model was optimized with a cross-entropy loss function using the Adam optimizer with a learning rate of $1e^{-4}$ and mini-batches of size 64. For CLEVR, we use the Stacked Attention Network [29] with

⁴We make the code and data available in supplementary material.

⁵We were unable to replicate the results of [24] on CLEVR dataset. This is why we used another baseline (Stacked Attention Network) and show how our method improves on that baseline. The primary reason being the original network was trained by authors on 10 parallel GPUs on 640 batch size. This was not feasible to replicate in lab setting. Based on our experiments, the best accuracy obtained by the relational reasoning network is 68% with a batch-size of 640 on a single-GPU worker, after running for 600 epochs over the dataset.

⁶A four-layer MLP consisting of 2000 units per layer with ReLU non-linearities is used for g_θ ; and a four-layer MLP consisting of 2000, 1000, 500, and 100 units with ReLU non-linearities used for f_ϕ .

the similar convolutional network and LSTM as above. We get similar results with VGG-16 as the convolutional network. Instead of the RN layer, we pass the two outputs through two levels of stacked attention, followed by a fully-connected layer. On top of this basic architecture, we define the student and teacher networks. The student network uses the same architecture as the baseline. We propose two variations of the teacher network, and we empirically show how these proposed changes improve upon the performance of the baseline network.

4.2. External Mask Prediction

In this setting of the experiment, the ground-truth mask, as calculated in 3.2.2, is element-wise multiplied to the image and then the image is passed through the convolutional network. We experiment with both sequential and iterative knowledge distillation. In the sequential setting, we first train the teacher network for 100 epochs with random embedding size of 32, batch size as 64, learning rate 0.0001. In the previous attempts to use distillation in natural language processing ([10, 13]), the optimal value of π has been reported as $\min(0.9, 1 - 0.9^t)$ or 0.9^t . Intuitively, either at the early or at the latter stages, the student almost completely *trusts* the teacher. However, our experiments show different results. For the student network, we employ a hyperparameter search on the value of imitation parameter π and use two settings, where π is fixed throughout the training and in the second setting, π is varied using $\min(\pi, 1 - \pi^t)$. We vary the loss ℓ_2 among cross entropy and euclidean norm. The results of the hyperparameter optimization experiment is depicted in Fig. 5. From this experiment, it can be observed that varying π over epochs gives better results than using a fixed π value for training the student. We observe a sharp increase in accuracy using the π value 0.575. This result is more consistent with the parameter value chosen by the authors in [30]. We also experiment by varying the word embedding (50-dimensional glove embedding and 32-dimensional word embedding) and learning rate. For sequential knowledge distillation, we get the best results with glove embedding and learning rate as $1e^{-4}$. However, we get huge improvements by using iterative knowledge distillation, where in each alternate epoch the student learns from the teacher and the groundtruth data; and the teacher learns from its original loss function and the student’s soft prediction (similar to Eqn. 5). Both weighted loss functions use the imitation parameter 0.9 (which remains fixed during training). We show the gradual learning of the teacher and the student till 800 epochs in Fig. 4 and compare it with the RN baseline. We observe that: 1) the External Mask-augmented Teacher network converges faster than the baseline and 2) the Student network outperforms the baseline after 650 epochs of training.

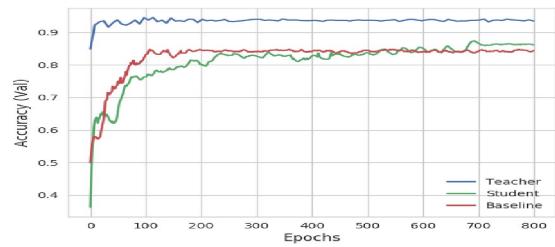


Figure 4: Validation accuracy after each epoch for teacher and student networks for iterative knowledge distillation on Sort-of-Clevr dataset and compare with the baseline.

4.3. Larger Model with Attention

In this framework, we investigate whether the mask can be learnt inside the network with attention mechanism. We train the teacher network for 200 epochs with glove vectors of size 50, batch size as 64, learning rate as 0.0001. We have employed a hyperparameter search over learning rate, embedding type, and learning rate decay, and found that the above configuration produces best results. For the student network, we employed a similar hyperparameter search on the value of imitation parameter π and use two settings, where π is fixed throughout the training and in the second setting, π is varied using $\min(\pi, 1 - \pi^t)$. We also vary the learning rate and the type of embedding (random with size 32 or glove vectors of size 50). The effect of the hyperparameter search is plotted in Fig. 6. We have experimented with iterative knowledge distillation and the best accuracy obtained for the teacher and the student networks are similar to that of sequential setting. The best test accuracies of the student network, the teacher with larger model and the baselines are provided in Table 1.

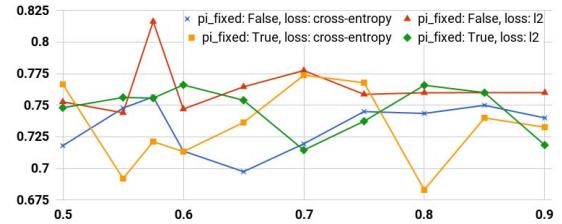


Figure 5: External Mask Prediction: Test Accuracy for different hyperparameter combination to obtain the best imitation parameter (π) for student in sequential distillation.

4.4. Analysis

The reported baseline accuracy on Sort-of-Clevr by [24] is 94% for both relational and non-relational questions. However, we use LSTMs to embed the natural language questions. Our implementation of the baseline achieves an

	Baseline	External Mask		In-Network Mask		Performance Boost Over Baseline (Δ)	
		Teacher	Student	Teacher	Student	Teacher	Student
Sort-of-Clevr	82% ([24])	95.7%	88.2%	87.5%	82.8%	13.7%	6.2%
CLEVR	53% ([29])	58%	55%	-	-	5%	2%

Table 1: Test set accuracies of different architectures for the Sort-of-Clevr (with natural language questions) and CLEVR dataset. For CLEVR, we used the Stacked Attention Network (SAN) [29] as baseline and only conducted the external-mask setting experiment as it already calculates in-network attention. Our re-implementation of SAN achieves 53% accuracy on CLEVR. Accuracy reported by [24] on SAN is 61%. The reported best accuracy for Sort-of-Clevr and CLEVR are 94% (one-hot questions [24]) and 97.8% ([21]).

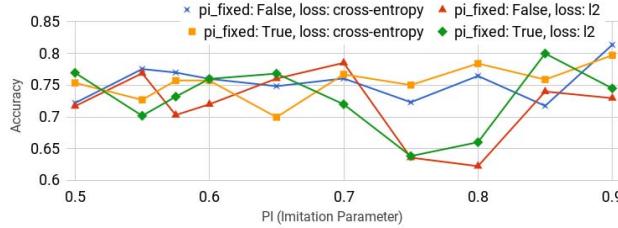


Figure 6: Model with Attention Mask: Test accuracy for the student network for different hyperparameter combination to obtain the best imitation parameter (π). We get the best validation accuracy using the π as 0.9, ℓ_2 as cross entropy loss and varying π over epochs.

overall test accuracy of 89% with one-hot question representation and 82% with LSTM embedding of the question. Addition of the pre-processed mask provides an increase in test accuracy to **95.7%**. In contrast, the teacher model with attention mask achieves **87.5%**. This is expected as the mask on the image simplifies the task by eliminating irrelevant region of image with respect to the question.

Student Learning: One may argue that adding such additional information to a model can be an unfair comparison. However, in this work, our main aim is to integrate additional knowledge (when it is available) with existing neural network architectures and demonstrate the benefits that such knowledge can provide. We experiment with the knowledge distillation paradigm to distill knowledge to a student. Extracted knowledge can be noisy, imperfect and often costly at test time. The distillation paradigm helps in this regard as the student network can choose to learn from the ground-truth data (putting less weight on teacher's predictions) during the training phase and doesn't require the additional knowledge during test time. For Sort-of-Clevr, we see an accuracy of **88.2%** achieved by the student network (in external mask setting), whereas for CLEVR the distillation effort increases the accuracy over the baseline method by 2%. Lastly, we show some qualitative examples of student network's output on the Sort-of-Clevr dataset (Fig. 7). The qualitative results indicate that our method can handle counting, spatial relationships well, but fails mostly on cases relating to shapes. This observation cou-

pled with improvement in generalization validates that the spatial knowledge has a significant role in our method.

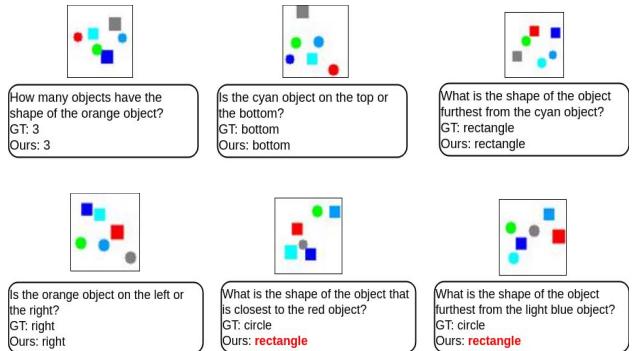


Figure 7: Some example images, questions and answers from the synthetically generated Sort-of-Clevr dataset. Red-colored answers indicate failure cases.

5. Conclusion

There has been a significant increase in attempts to integrate background knowledge (linguistic knowledge [30] or commonsense rules [10]) with state-of-the-art Neural Architectures in Computer Vision and Natural Language Processing applications. In this work, we attempt to integrate additional information in the form of spatial knowledge with existing neural networks to aid Visual Reasoning. The spatial knowledge is obtained by reasoning on the natural language question and additional scene information using the Probabilistic Soft Logic inference mechanism. We show that such information can be encoded using a mask over the image and integrated with neural networks using knowledge distillation. Such a procedure shows significant improvement on the accuracy over the baseline network.

6. Acknowledgements

The support of the National Science Foundation under the Robust Intelligence Program (1816039 and 1750082), and a gift from Verisk AI are gratefully acknowledged. We also acknowledge NVIDIA for the donation of GPUs.

References

[1] S. Antol, A. Agrawal, J. Lu, M. Mitchell, D. Batra, C. Lawrence Zitnick, and D. Parikh. Vqa: Visual question answering. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 2425–2433, 2015.

[2] S. H. Bach, M. Broecheler, B. Huang, and L. Getoor. Hinge-loss markov random fields and probabilistic soft logic. *Journal of Machine Learning Research (JMLR)*, 2017. To appear.

[3] D. Bahdanau, K. Cho, and Y. Bengio. Neural machine translation by jointly learning to align and translate. *arXiv preprint arXiv:1409.0473*, 2014.

[4] C. Baral. *Knowledge representation, reasoning and declarative problem solving*. Cambridge university press, 2003.

[5] A. G. Cohn and J. Renz. Chapter 13 qualitative spatial representation and reasoning. In F. van Harmelen, V. Lifschitz, and B. Porter, editors, *Handbook of Knowledge Representation*, volume 3 of *Foundations of Artificial Intelligence*, pages 551 – 596. Elsevier, 2008.

[6] L. De Raedt, A. Kimmig, and H. Toivonen. Problog: A probabilistic prolog and its application in link discovery. In *Proceedings of the 20th International Joint Conference on Artificial Intelligence, IJCAI'07*, pages 2468–2473, San Francisco, CA, USA, 2007. Morgan Kaufmann Publishers Inc.

[7] D. Elliott and F. Keller. Image description using visual dependency representations. In *Proceedings of the 2013 Conference on Empirical Methods in Natural Language Processing*, pages 1292–1302, 2013.

[8] G. Hinton, O. Vinyals, and J. Dean. Distilling the Knowledge in a Neural Network. <http://arxiv.org/pdf/1503.02531v1.pdf>, 2015.

[9] R. Hu, J. Andreas, M. Rohrbach, T. Darrell, and K. Saenko. Learning to reason: End-to-end module networks for visual question answering. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, 2017.

[10] Z. Hu, X. Ma, Z. Liu, E. Hovy, and E. Xing. Harnessing deep neural networks with logic rules. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2410–2420, Berlin, Germany, August 2016. Association for Computational Linguistics.

[11] Z. Hu, Z. Yang, R. Salakhutdinov, and E. Xing. Deep neural networks with massive learned knowledge. In *Proceedings of the 2016 Conference on EMNLP*, pages 1670–1679, Austin, Texas, November 2016. ACL.

[12] J. Johnson, B. Hariharan, L. van der Maaten, L. Fei-Fei, C. L. Zitnick, and R. Girshick. Clevr: A diagnostic dataset for compositional language and elementary visual reasoning. *arXiv preprint arXiv:1612.06890*, 2016.

[13] Y. Kim and A. M. Rush. Sequence-level knowledge distillation. *arXiv preprint arXiv:1606.07947*, 2016.

[14] G. Klir and B. Yuan. Fuzzy sets and fuzzy logic: theory and applications. 1995.

[15] B. M. Lake, T. D. Ullman, J. B. Tenenbaum, and S. J. Gershman. Building machines that learn and think like people. *Behavioral and Brain Sciences*, pages 1–101, 2016.

[16] Y. LeCun. A path to ai. <https://futureoflife.org/wp-content/uploads/2017/01/Yann-LeCun.pdf>, 2017.

[17] B. London, S. Khamis, S. Bach, B. Huang, L. Getoor, and L. Davis. Collective activity detection using hinge-loss markov random fields. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops*, pages 566–571, 2013.

[18] D. Lopez-Paz, L. Bottou, B. Schlkopf, and V. Vapnik. Unifying distillation and privileged information, 2015. [arxiv:1511.03643](https://arxiv.org/abs/1511.03643).

[19] J. Mun, M. Cho, and B. Han. Text-guided attention model for image captioning. In *AAAI*, pages 4233–4239, 2017.

[20] J. Pennington, R. Socher, and C. D. Manning. Glove: Global vectors for word representation. In *Empirical Methods in Natural Language Processing (EMNLP)*, pages 1532–1543, 2014.

[21] E. Perez, F. Strub, H. De Vries, V. Dumoulin, and A. Courville. Film: Visual reasoning with a general conditioning layer. *arXiv preprint arXiv:1709.07871*, 2017.

[22] D. A. Randell, Z. Cui, and A. G. Cohn. A spatial logic based on regions and connection. In *Proceedings 3rd International Conference ON Knowledge Representation And Reasoning*, 1992.

[23] M. Richardson and P. Domingos. Markov logic networks. *Mach. Learn.*, 62(1-2):107–136, Feb. 2006.

[24] A. Santoro, D. Raposo, D. G. Barrett, M. Malinowski, R. Pascanu, P. Battaglia, and T. Lillicrap. A simple neural network module for relational reasoning. *arXiv preprint arXiv:1706.01427*, 2017.

[25] V. Vapnik and R. Izmailov. Learning using privileged information: similarity control and knowledge transfer. *Journal of machine learning research*, 16:55, 2015.

[26] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. Kaiser, and I. Polosukhin. Attention is all you need. *arXiv preprint arXiv:1706.03762*, 2017.

[27] P. Wang, Q. Wu, C. Shen, A. Dick, and A. van den Hengel. Fvqa: fact-based visual question answering. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2017.

[28] K. Xu, J. Ba, R. Kiros, K. Cho, A. Courville, R. Salakhudinov, R. Zemel, and Y. Bengio. Show, attend and tell: Neural image caption generation with visual attention. In *ICML*, pages 2048–2057, 2015.

[29] Z. Yang, X. He, J. Gao, L. Deng, and A. Smola. Stacked attention networks for image question answering. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 21–29, 2016.

[30] R. Yu, A. Li, V. I. Morariu, and L. S. Davis. Visual relationship detection with internal and external linguistic knowledge distillation. *IEEE International Conference on Computer Vision (ICCV)*, 2017.

[31] S. Zheng, S. Jayasumana, B. Romera-Paredes, V. Vineet, Z. Su, D. Du, C. Huang, and P. H. S. Torr. Conditional random fields as recurrent neural networks. In *Proceedings of the 2015 IEEE International Conference on Computer Vision (ICCV)*, ICCV '15, pages 1529–1537, Washington, DC, USA, 2015. IEEE Computer Society.