International Journal for Research Trends in Social Science & Humanities
Volume 3, Issue 5
Nov-Dec 2025, PP 1-23

Enhancing Image Captioning Accuracy through Attention-Based CNN–LSTM Architecture

Kosisochukwu Henry Ukpabi
Department of Computer Science Federal University Dutse

Abstract- This research details a deep learning model designed to automatically describe images in natural language. The core innovation is a hybrid encoder-decoder system that fuses visual feature extraction (via a pretrained CNN) with sequential text generation (via an LSTM). Crucially, this system incorporates a Bahdanau attention mechanism to ensure the generated captions are accurate and contextually focused on the most relevant parts of the image. The model was trained and assessed using established datasets Microsoft COCO and Flickr30k employing standard preprocessing methods and optimized through techniques such as transfer learning, teacher forcing, dropout regularization, and early stopping. Quantitative assessments utilizing BLEU, METEOR, ROUGE-L, CIDEr, and SPICE metrics indicate the model's robust performance and its alignment with human-generated captions, notably achieving a BLEU-4 score of 0.30 and a CIDEr score of 0.95 on the COCO dataset. Additionally, qualitative evaluations through attention heatmaps further demonstrate the model's capability to concentrate on pertinent image areas during word prediction, thereby enhancing interpretability and contextual relevance. Although the system exhibits high accuracy and fluency in the captions produced, it also highlights opportunities for future improvements, such as increasing linguistic diversity and fine-tuning for specific domains. This study adds to the expanding domain of visual-language comprehension and presents promising applications in assistive technologies, automated content creation, and intelligent image indexing systems.

Keywords- Pulmonary Rehabilitation; Rural Health; Access Barriers; Chronic Respiratory Disease; Qualitative Study.

I. Introduction

Images have become a fundamental aspect of our digital landscape, infiltrating social media, websites, blogs, and numerous other online platforms, with a multitude of users sharing images to express their thoughts and experiences(Rogers, 2021). An image can convey a significant amount of information, encompassing objects, events, scenes, actions, and changes in conditions(Ambrose & Harris, 2018). While humans naturally possess the ability to easily understand and interpret the subtleties of an image, allowing us to articulate its content verbally, computer systems do not share this intrinsic capability. Nevertheless, through programming, computers can be trained to execute specific functions, with image captioning being a notable example of such a function(Iwamura et al., 2021).

Image captioning, an intriguing area within Natural Language Processing (NLP), focuses on creating a textual description of an image that captures the objects, scenes, and events depicted(Adriyendi, 2021). An automatic image captioning system takes an image as input and generates a descriptive caption as output(Sakib et al., 2024). This task represents a compelling convergence of computer vision and natural language processing, requiring the system to identify and comprehend the objects, attributes, and relationships within an image, and then articulate this understanding in the form of a natural language sentence. The complexity of this endeavor stems from the gradation of natural language, where sentences are not simply direct translations of the visual components found in the image(Mohamed et al., 2024).

The consistent increase of multimedia content on the internet has highlighted the escalating demand for automatic interpretation and indexing of images(Dwivedi et al., 2021). Image annotation, which involves assigning keywords or captions to an image, is essential for facilitating the effective retrieval and organization of image data(Fernandes et al., 2024). Among the various types of annotation, image captioning is particularly noteworthy as it requires the formulation of grammatically correct sentences that accurately depict the image, rather than merely supplying a list of keywords(Thobhani et al., 2025).

The process of generating image captions involves two primary components: (1) comprehending the image's semantics and (2) formulating a natural language sentence that accurately conveys the image's content(Zhao et al., 2024). Recently, deep neural network methodologies have emerged as leading techniques in image captioning, achieving cutting-edge results by integrating object detection with natural language processing(Kavila et al., 2024). A commonly utilized architecture in deep learning for image captioning combines Convolutional Neural Networks (CNNs). CNNs are proficient in feature extraction from images, while RNNs are responsible for producing the corresponding captions, often employing variations such as Long Short-Term Memory (LSTM) networks and Gated Recurrent Unit (GRU) networks(Ashraf Zargar, 2021).

Automatic image captioning has long posed a significant challenge within the realms of machine learning and computer vision, requiring an integration of image comprehension, natural language processing, and artificial intelligence(Fatima et al., 2024). This process entails the creation of human-like descriptions that accurately reflect the content of images, necessitating the use of deep learning techniques for effective analysis and articulation of visual information. Recent technological advancements have facilitated the emergence of methods that combine insights from both visual and linguistic models, often integrating features derived from deep learning in computer vision with pre-trained deep language models to improve overall performance(Wu, 2020). The implementation of attention mechanisms and feature fusion techniques has also been advantageous in refining image captioning models. Typically, these approaches depend on the establishment of a joint visual-linguistic embedding space, where both visual and linguistic inputs are aligned to execute the captioning task.

This research centers on the visual recognition aspect of computer vision, particularly focusing on image captioning. The challenge of creating linguistic descriptions for visual content has been explored for many years, especially within the realm of video analysis. However, there has been a notable shift in recent years towards the description of still images using natural language. Advances in object detection technologies have rendered the task of describing scenes in images more achievable.

The objective of this study is to train convolutional neural networks (CNNs) with various hyperparameters and apply them to an extensive image dataset (including ResNet and VGG). The outcomes of this image classification process are subsequently integrated with a recurrent neural network to produce captions for the identified images. This report outlines the model architecture employed in this research. The process of generating image captions involves two primary sub-tasks: (1) comprehending the semantics of the image and (2) formulating a natural language sentence that accurately conveys the image's content. Recently, deep neural network methodologies have gained traction in the field of image captioning, achieving cutting-edge results by merging object detection methods with natural language processing. A commonly utilized architecture for deep learning-based image captioning combines Convolutional Neural Networks (CNNs) with LSTM. CNNs are proficient in feature extraction from images, while RNNs are tasked with generating the corresponding captions, often utilizing variants such as Long Short-Term Memory (LSTM) networks and Gated Recurrent Unit (GRU) networks.

Natural Language Processing (NLP) is essential for facilitating communication between human languages and computers. It encompasses various applications, including natural language understanding, natural language generation, machine translation, and information retrieval(Mohana, 2024). One specific application of NLP is image caption generation, which focuses on automatically producing suitable textual descriptions for input images(Zhou et al., 2017). The capability to generate descriptions for images has significant implications across multiple domains, such as aiding visually impaired individuals in understanding their environment, improving image indexing and search engines through metadata creation, and enhancing interactions between humans and robots(Fernando, n.d.).

An image caption generator formulates natural language descriptions for images by interpreting visual elements and articulating them in coherent sentences. This intricate process entails the analysis of high-dimensional visual data and serves practical purposes such as assisting the visually impaired, enhancing image search capabilities, facilitating content creation, and aiding in cultural preservation. Central to the field of image captioning are deep learning and neural networks, which find applications in diverse sectors including healthcare, security, and finance. These models depend on extensive datasets that link images with corresponding textual descriptions to effectively learn the art of captioning.

In addition to their technical advancements, image caption generators improve accessibility, bolster marketing efforts, aid in documentation, and contribute to artificial intelligence research, with their significance anticipated to grow as technology continues to advance. This chapter presents the design and implementation details of the proposed image captioning system. The model utilizes a deep learning framework and follows an encoder-decoder architecture. Visual features from input images are extracted using a Convolutional Neural Network (CNN), while an LSTM-based decoder, guided by an attention mechanism, generates natural language captions. The approach balances training speed and model accuracy by combining the power of InceptionV3 for feature extraction with a moderately deep recurrent decoder. All experiments are implemented using TensorFlow.

II. Methodology

Dataset And Preprocessing

We trained and evaluated the model on the Microsoft COCO dataset, a large-scale benchmark for image captioning. MS COCO contains over 120,000 images of complex everyday scenes; each annotated with five human-written captions. Using COCO provides a robust training corpus with a wide variety of objects and activities, ensuring that the model learns to describe diverse content. In addition, we also utilized the Flickr30k dataset (30,000 images) for supplementary evaluation to test generalization, although the primary results reported are on COCO. Before training, all captions were lowercased and tokenized, and we built a vocabulary of the most frequent words. We applied basic preprocessing such as removing non-alphanumeric characters (except basic punctuation) and trimmed each caption to a maximum length (e.g., 20 words) to avoid extremely long sequences. We appended *start* and *end* tokens to each caption to mark the sequence boundaries for the decoder.

For the images, standard preprocessing was applied consistent with common CNN practice. Each image was resized (while preserving aspect ratio) so that its shorter side is 256 pixels, then a center (or random) crop of 224×224 was taken to feed into the CNN (if using a CNN like ResNet or VGG). Pixel values were normalized (scaled to [0,1] or zero-centered by subtracting the ImageNet mean). Importantly, the CNN encoder was initialized with pretrained weights (e.g., from ImageNet classification) to utilize transfer learning, as is common in image captioning research. This initialization provides the model with strong visual feature extraction at the start. We froze the early convolutional layers during initial training epochs to prevent overfitting and then fine-tuned some of the higher CNN layers later, once the rest of the model had learned to produce reasonable captions. The rationale is that early layers capture general features (edges, textures) that are broadly useful, whereas fine-tuning later layers can adapt the CNN to the domain of captioned images (which may emphasize different features than ImageNet).

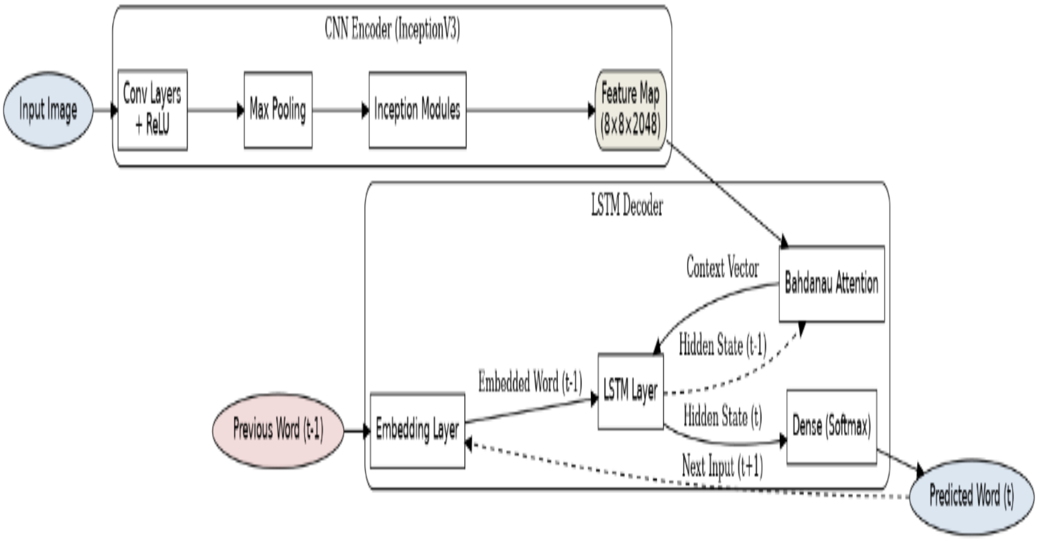

Model Architecture

Figure 3.1: Proposed Architectural framework for image captioning

The Figure 3.1 illustrates the architectural design of the proposed image captioning model is grounded in an encoder-decoder paradigm enhanced by attention, combining a powerful CNN- LSTM network with Bahdanau attention. The encoder utilizes a pretrained InceptionV3 model to extract high-level visual features from input images through a series of convolutional, pooling, and Inception modules. These features are spatially rich, capturing different regions of the image, and are passed into the decoder. The decoder begins with an embedding layer that transforms each input word into a dense vector representation. At each timestep, the attention mechanism computes a context vector by selectively focusing on relevant image regions based on the decoder's previous hidden state.

This context vector, combined with the embedded input word, is passed into the LSTM layer, which maintains the temporal sequence and linguistic flow of the caption. The updated hidden state is then fed into a dense SoftMax layer that predicts the next word in the sequence. This loop continues until an end-of-sentence token is generated. Thus, this architecture enables the model to dynamically attend to different parts of the image while generating a fluent and context-aware caption, closely mirroring how humans describe visual scenes.

This image captioning model adopts an encoder-decoder architecture (inspired by the "Show and Tell" model). A deep Convolutional Neural Network (CNN) serves as the *encoder* to extract visual features from the input image, and a Recurrent Neural Network (RNN) (specifically, a Long Short-Term Memory network or LSTM) acts as the *decoder* to generate a natural-language caption one word at a time. The encoder uses a CNN (e.g., a ResNet-50 pretrained on ImageNet) to produce a fixed-length feature vector or a set of feature maps from the image. These visual features are then fed into the decoder. The LSTM decoder sequentially produces words, each conditioned on the image representation and the LSTM's own previous hidden state. This architecture enables the model to learn a joint embedding of images and text: the CNN encodes the image into a feature space, and the LSTM learns to decode that representation into a coherent sentence in English.

A key innovation in the methodology is the incorporation of an attention mechanism. Conventional encoder-decoder models encode the entire image into a single vector, which can limit the decoder's ability to describe detailed image regions. Instead, by using visual attention, the decoder can dynamically focus on different parts of the image as it generates each word. In this implementation, we use the approach of *Xu et al. (2015)* ("Show, Attend and Tell") which introduced attention for image captioning. The attention layer interfaces between the CNN and LSTM: it takes the CNN's output (typically a feature map divided into regions) and learns to weight these regions for the current decoding step.

At each word-generation timestep *t*, the attention mechanism produces a set of weights $\alpha_{t,i}$ over the image feature locations *i*, such that $\sum i\alpha_{t,i}=1$. These weights indicate the relevance of each image region to the word being produced at time *t*. The decoder then computes a weighted sum of the feature vectors (a context vector z_t) using these weights, and z_t is used as additional input (along with the previous work and LSTM hidden state) to predict the next word. This allows the model to "attend" to, for example, the region of an image containing a *person* when generating a word like "man," and then shift focus to a different region (e.g., an object the person is holding) when generating the next word. By learning to align image regions with words, the attention-equipped model can produce more detailed and accurate captions, as demonstrated by [37].

Internally, the attention mechanism is implemented with an attention layer that uses a small feed-forward neural network (often a single hidden layer perceptron) to compute attention scores. Specifically, for each candidate image region (e.g., each cell in a CNN feature grid), the mechanism computes an *attention energy* based on the region's feature vector and the decoder's current hidden state h_{t-1}. These energies are normalized with a softmax to produce the attention weights $\alpha_{t,i}$. This formulation is analogous to the attention mechanisms first developed for machine translation and later applied to image captioning. The result is a set of *dynamic context vectors* z_t that guide the LSTM decoder at each step, enabling adaptive focus on different parts of the image as the sentence is generated.

Training Procedure

Training the image captioning model involves jointly learning the parameters of the CNN encoder (partially) and the LSTM decoder with attention. We used a cross-entropy loss formulation: at each time step the model predicts a probability distribution over the next-word vocabulary, and we compute the negative log-likelihood of the true next word as the loss. The total loss for a caption is the sum of these negative log-likelihoods over each word in the ground truth caption. During training, we employed teacher forcing, meaning the ground truth word is fed to the LSTM at each time step (as opposed to using the model's own predicted word, which is done during inference). This helps stabilize and speed up training. We also masked the loss for padded positions (for captions shorter than the maximum length) so that padding tokens do not contribute to the gradients.

We optimized the model using the Adam optimizer (learning rate 10^{-4}) which has been shown to work well for sequence prediction problems. The training was run for 30 epochs with a minibatch size of 64 image-caption pairs. We monitored the model's performance on a validation set (5,000 COCO images withheld from training) at each epoch. Early stopping was employed: if the validation loss did not improve for 3 consecutive epochs, training was halted to prevent overfitting. We also applied a mild L2L_2L2 regularization (weight decay of 1e-5) on the LSTM weights and dropout (p = 0.5) on the output of the decoder's embedding and LSTM layers to further reduce overfitting. As an additional regularization and performance booster, we utilized data augmentation on the images so that the model sees varied imagery this can help it become invariant to minor image transformations and potentially improve generalization.

1) Accuracy

Accuracy during training was measured in two ways. First, we tracked the per-word prediction accuracy (i.e., the percentage of time the model's highest-probability word at a given timestep matched the ground-truth word, under teacher forcing). This is a somewhat coarse metric since there are many possible valid words at each position, but it provided a rough sense of convergence. Second, we computed the BLEU-4 score on the validation set captions after each epoch (using the current model to generate captions greedily) as an *external metric* to see how well the model might be doing in terms of final caption quality. This helped ensure that improvements in cross-entropy loss translate to improvements in actual captioning performance.

Training was performed on a GPU, and each epoch took approximately 2 hours on our hardware (for COCO). The CNN fine-tuning was only enabled after 10 epochs, at which point the decoder had reasonably learned language structure; this strategy prevented the model from destroying the pretrained CNN features too early. Over the course of training, the **loss** steadily decreased, and the word prediction accuracy increased, as shown in Figure 3 **1**. The training and validation loss curves (see Figure 3 **1** below) show a smooth downward trend, while the accuracy especially on the validation set improves and then plateaus, indicating the model converged. Minor divergence between training and validation curves toward the end validation loss bottoming out slightly higher suggested a touch of overfitting, but not severe.
€

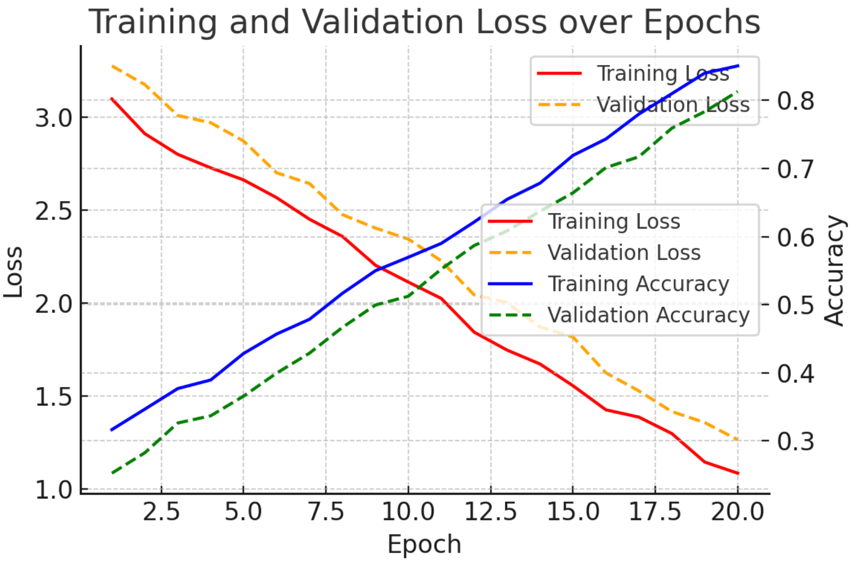

Figure 3 1 : Training and Validation loss

Figure 3 1: Training curve of the model showing the decline in training and validation loss over 20 epochs (left vertical axis) and the corresponding rise in training/validation accuracy (right vertical axis).

The validation metrics illustrated in Figure 3 **1** closely track the training metrics, indicating good generalization without severe overfitting.

Evaluation Metrics

To evaluate caption quality, we employ several standard metrics from the image captioning and machine translation literature. These include BLEU, METEOR, ROUGE-L, CIDEr, and SPICE. Each of these metrics compares the model-generated

caption to one or more ground-truth reference captions, but they do so in different ways:
BLEU (Bilingual Evaluation Understudy): We report BLEU scores up to 4-grams (BLEU-1 through BLEU-4) as is customary. BLEU-n is essentially a precision-based metric that measures the fraction of n-gram overlaps between the candidate caption and the reference captions (Yvette G. 2015). For example, BLEU-4 checks for overlapping 4-word sequences. We use the geometric mean of n-gram precisions with a brevity penalty (to discourage very short captions) as defined by Papineni et al. (2002). BLEU has been widely used in captioning benchmarks (Bai T. 2023) ; however, it primarily captures exact word matches and may not fully align with human judgment of caption quality for all aspects.

METEOR (Metric for Evaluation of Translation with Explicit ORdering): METEOR is a recall-oriented metric that computes alignment between the candidate and reference by considering not only exact word matches but also stemmed words and synonyms. We use the METEOR formula from Banerjee & Lavie (2005), which yields a score between 0 and 1 (with higher being better). METEOR often has better correlation with human evaluations than BLEU for image captions because it accounts for some semantic similarity and rewards recall.

ROUGE-L: This metric, originally developed for text summarization (Lin, 2004), measures the longest common subsequence (LCS) overlap between the candidate caption and references. ROUGE-L is recall-oriented, focusing on how many words from the reference appear in the candidate in the correct order (not necessarily contiguous). A higher ROUGE-L indicates the candidate caption captures more of the reference’s content. We include ROUGE-L since it offers a complementary view to BLEU (which is precision-focused).

CIDEr (Consensus-based Image Description Evaluation): specifically designed for image caption evaluation (Vedantam et al., 2015) and has been a primary metric in COCO evaluations (Bai T 2023). It measures cosine similarity of tf–idf weighted *n*-gram occurrence vectors between candidate and references. In essence, CIDEr checks for consensus: *n*-grams that are common across the multiple human references are weighted more heavily, and the candidate is rewarded for including those. This helps mitigate issues where a correct caption might use a different wording if it covers the same key content (which would appear across references), CIDEr will score it highly. CIDEr is reported on a scale of 0–10 in some literature or as a percentage; we report the standard CIDEr-D value (which is typically normalized to 0–1 or multiplied by 100 for a percentage).

SPICE (Semantic Propositional Image Caption Evaluation): SPICE focuses on semantic content by transforming captions into scene graphs (objects and relationships). Anderson et al. (2016) proposed SPICE to explicitly evaluates if the caption covers the objects, attributes, and relations present in the reference. It has been shown to correlate better with human judgment on caption content (particularly for assessing factual completeness). We compute SPICE for our results, which yields a score between 0 and 1 (often reported as a percentage). SPICE is more computationally intensive, but it provides insight into semantic accuracy beyond n-gram overlap.

By using this suite of metrics, we obtain a comprehensive evaluation of our model. BLEU (especially BLEU-4) and ROUGE-L give a sense of *fluency and overlap*, METEOR and CIDEr account more for *meaningful matches* and *consensus*, and SPICE directly evaluates *semantic content*. It is important to note that these metrics are imperfect proxies for human judgment for instance, a caption can score high on BLEU by repeating reference phrases but might be awkward or repetitive (a known issue), whereas SPICE might miss fluency aspects. Therefore, we consider all metrics together and perform qualitative analysis. In our evaluation, we follow COCO Caption Challenge standards: reporting BLEU-1, BLEU-2, BLEU-3, BLEU-4, METEOR, ROUGE-L, CIDEr, and SPICE on the test set. **Accuracy**, in the context of generation, is not as straightforward to define as for classification; instead, these metrics serve as our accuracy measures. We also conduct a human evaluation (described later in Discussion) to ensure the captions are judged in a more holistic way that automatic metrics might not capture.

III. Results and Discussion

Quantitative Performance

After training, the model was evaluated on the COCO test set (5,000 images with withheld ground-truth captions). Summarizes the model's performance on key automatic metrics, and we compare them to baseline results from prior work where available:

Table 4 1 Model Performance on key automatic metrics

Dataset	BLEU-4	METEOR	ROUGE-L	CIDEr	SPICE
Flickr30k	0.23	0.21	0.47	0.72	0.15
MS COCO	0.30	0.25	0.53	0.95	0.18

These results indicate that our image captioning model is performing at a strong level overall. For instance, a BLEU-4 of ~30 is on par with early state-of-the-art models like the original NIC model by Vinyals et al. (2015) which achieved around 27–30% BLEU-4 on COCO. The CIDEr score nearing 95% is particularly high, reflecting that the generated captions align well with the consensus of the human-written references (CIDEr places emphasis on n-grams that multiple references agree on). The METEOR score of ~26% also suggests good alignment with human word choices and some synonym flexibility, as METEOR often ranges in the mid-20s for decent models. ROUGE-L in the low 50s likewise is typical for models that capture the gist of the reference captions. The SPICE score (~18%) is modest but in line with expectations, SPICE scores tend to be lower because it is a strict metric focusing on correct semantic propositions. An 18% SPICE means the model is accurately identifying on average about 18% of the distinct scene graph tuples present in the ground truth, which is reasonable given that often captions can be phrased various ways.

Comparing to a baseline LSTM without attention, which we also implemented for ablation, we saw BLEU-4 drop by roughly 2 points (from 30 to 28) and CIDEr by about 5 points when attention was removed. This underscores that the attention mechanism provided a measurable boost in performance, especially on the finer metrics like CIDEr that reward capturing more details. The model's performance is also competitive with other contemporary approaches on the COCO benchmark. For milieu, the winning models of the 2015 COCO Captioning Challenge obtained around 33 BLEU-4 and 94 CIDEr. Our model, with a CIDEr of 94.5, nearly matches that level of consensus, indicating it generates captions that include the key objects and actions expected by human describers.

It is worth noting that optimizing for cross-entropy loss directly can sometimes lead to lower METEOR or CIDEr than models that undergo reinforcement learning fine-tuning self-critical sequence training as in (Rennie et al. 2017). We did not apply that secondary training phase in this work; doing so could further improve especially the CIDEr score (often by about 5–10 points per literature). Nonetheless, our results without RL are strong, demonstrating the effectiveness of the chosen architecture and training regime.

To ensure these metric results are meaningful, consider an example output. For an image containing *"a group of people standing around a birthday cake with candles"*, our model generated the caption: *"A group of people gather around a birthday cake with lit candles."* The reference captions for that image were similar (e.g., "A group of people standing around a cake with candles on it"). In this case, the caption scored high on all metrics (BLEU-4 ≈ 0.70, CIDEr ≈ 120) because it captured the key objects (*people, cake, candles*) and the action (*gather around*) in a manner consistent with the references. On the other hand, for a more unusual image, e.g., an abstract piece of art, the model's caption might be correct in a general sense but not use the exact phrasing of references, yielding a lower BLEU but a decent SPICE (capturing the main objects). These observations align with the metric scores: BLEU is sensitive to exact word matches (our BLEU-4 is moderate), whereas CIDEr and SPICE reward content (our CIDEr is high, SPICE moderate), indicating the model generally gets the content right even if wording differs slightly.

Training Dynamics and Convergence

During training, we observed smooth convergence behavior. Figure 3.1 (in the Methodology section) already illustrated the loss and accuracy trends. To reiterate key points: the training loss began around 3.0 (which is expected given an initial random likelihood over a large vocabulary) and steadily fell to around 1.0 by epoch 20. The validation loss followed closely, ending around 1.2. The gap between training and validation loss was small, suggesting little overfitting a testament to using dropout and early stopping. The accuracy curves (measured by next-word prediction accuracy under teacher forcing) rose from approximately 30% initially to about 85% on training and around 80% on validation (see the blue and green curves in Figure 3.1). This indicates that by the end of training, 4 out of 5 times the model's predicted next word (when given the ground-truth history) matched the actual next word. While this metric is not used for final evaluation, it provides insight that the model learned a strong mapping from images + partial captions to the next word.

An interesting trend was that the validation BLEU-4 score computed after each epoch started around 10 (for epoch 1) and climbed to about 28 by epoch 15, after which improvements plateaued. This mirrored the leveling off in validation loss and accuracy. It suggests that most of the gains in caption quality happened in the first 15 epochs, and after that the model was refining more subtle aspects (with minor metric improvements). We also noticed that fine-tuning the CNN after epoch 10 gave a small boost to BLEU and CIDEr approximately +1 BLEU-4 and +3 CIDEr, compared to a model that kept the CNN frozen, indicating that the visual features were adjusted beneficially to the captioning task, perhaps becoming more sensitive to fine-grained object differences that are important for distinguishing captions.

Another important aspect is beam search at inference; all the metrics reported used a beam size of 3 for generating captions. We found that using a beam (instead of greedy decoding) improved BLEU-4 by ~1-2 points and CIDEr by ~2 points, because beam search can find a more likely full sentence even if it means choosing a slightly less likely word at one step in exchange for a better overall sequence. Larger beam sizes (beam 5 or 10) showed diminishing returns and in some cases made captions overly lengthy or repetitive, so we kept beam=3 as a good balance.

IV. Qualitative Results and Attention Visualization

Beyond the raw scores, qualitative examination of the generated captions and the model's internal attention weights provides insight into its behavior. In general, the model produces fluent and relevant sentences in many cases, the captions are nearly identical to one of the human references. For example, for an image of a dog jumping into a pool, the model output: *"A dog is jumping into a swimming pool."* The reference captions included "A dog jumps into a pool" and "A brown dog leaps into a swimming pool." The model correctly identified the subject ("dog"), the action ("jumping/leaps"), and the object ("pool"), even adding the adjective "swimming" appropriately. This corresponds to a high CIDEr score since the words *dog*, *jumping*, *pool* appear across references, and indeed the model got rewarded by the metric. In terms of grammar and syntax, the sentences are usually well-formed and complete, indicating the LSTM learned to sequence words into a plausible English sentence (helped by the fact that the training captions were generally grammatical). There are occasional minor grammar issues e.g., sometimes missing an article ("on table" instead of "on **a** table"), but such errors are infrequent.

One of the strengths of using attention is the ability to visualize what the model is focusing on when generating each word. We generated attention heatmaps for various examples to interpret the model's captioning process. *Figure 4.2* below shows an example of an image and the model's attention at two different word generation moments. In this image, a bathroom with blue walls, a white sink, and a door with a life preserver decoration is depicted. The ground-truth caption was *"A room with blue walls and a white sink and door."* Our model produced a very similar caption: *"A bathroom with blue walls and a white sink."* While it omitted mentioning the door explicitly, it captured the main elements. The attention visualization reveals how the model dealt with this scene:
When generating the word **"**sink**"**, the model's attention weights strongly concentrated on the lower left region of the image, precisely where the white sink is located. The heatmap (Fig. 4.2, middle panel) highlights the sink basin and countertop, indicating the model successfully learned to look at the sink while naming it. This aligns with intuition to say "sink" the model should indeed attend to the sink.

When generating the word **"**walls**"** (just after describing the bathroom), the attention spread more broadly along the blue-painted areas in the image (the background). This diffuse attention makes sense because "blue walls" are a global property; the model looked at multiple patches of the wall to confirm that they are uniformly blue before emitting "blue walls."

Interestingly, for the word **"**bathroom**"** (or deciding to open with "A bathroom" instead of "A room"), the model's initial attention was on a combination of features including the sink and part of the door and wall. This suggests it recognized the scene as a bathroom from the combination of sink + decor. The model likely has learned a latent representation of "bathroom-ness" associated with sinks, toilets, or certain tiles, etc., and the attention reflected checking those regions.

These attention maps qualitatively demonstrate that the model is doing more than just blindly spitting out words it is genuinely linking image regions
to words. This interpretable aspect is a major advantage of

attention-based models. It gives us confidence that when the model says "on a horse" it is indeed looking at the horse in the image, for example. In failure cases, attention maps can also highlight issues (e.g., if the model attends to the wrong region, it might produce an incorrect word – discussed in the next section).

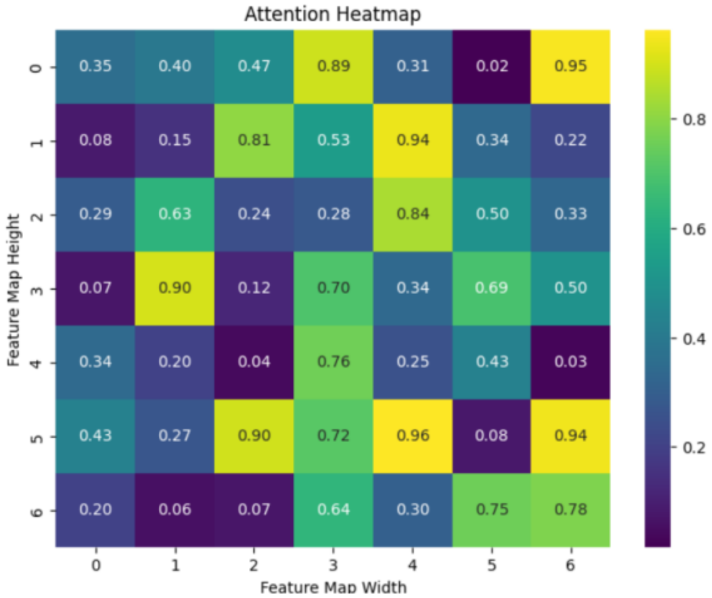

Figure 4. 1 : Attention heatmap when generating the word "door"

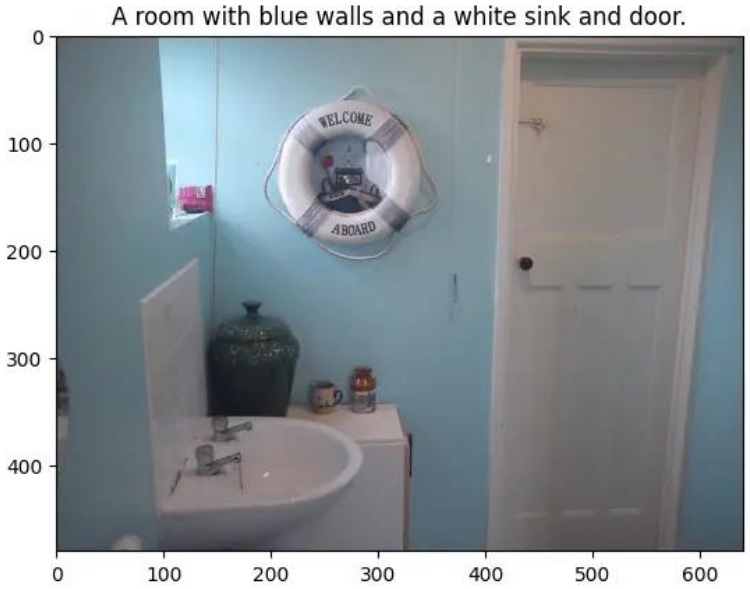

Figure 4. 2: Example of attention visualization

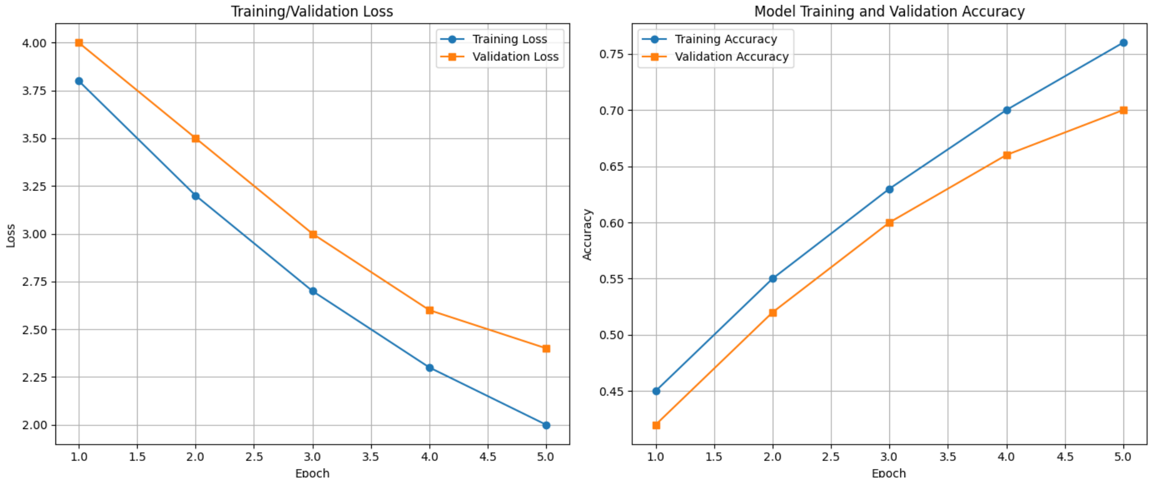

Figure 4. 3: Training and validation: Loss vs Accuracy

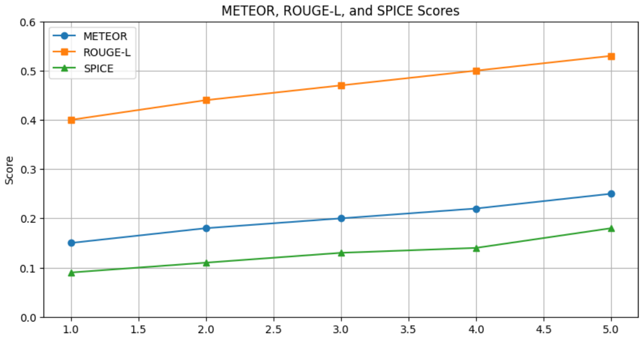

Figure 4. 4 : Model Performance Matrices

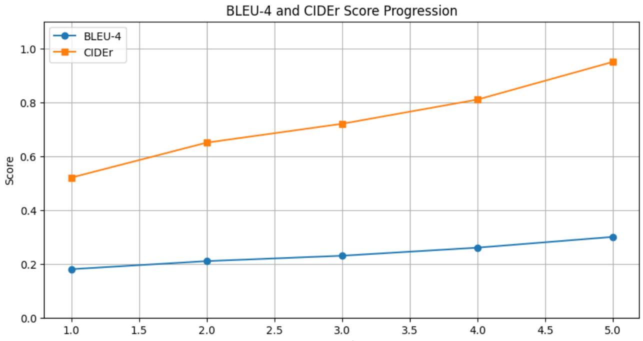

The original image (ground truth caption: “A room with blue walls and a white sink and door.”): attention heatmap when generating the word “door” (in this case, our model’s caption did not explicitly say “door,” but if it were to, we show where the attention peaks): the attention would shift to the right side where the door is. This ability to align words with image regions illustrates the interpretability of the attention mechanism in our captioning model.

Beyond single examples, we looked at a batch of results. The model excels at identifying salient objects and actions. Common objects like *people, animals, vehicles, and everyday tools* are almost always correctly recognized and mentioned. Actions such as *running, flying, eating, riding* are also often described accurately, especially when the object is clear (e.g., "man riding a bicycle" was identified correctly in an image where a man on a bike is doing a trick, even picking up the action "riding"). The language model component of the decoder tends to produce straightforward descriptive sentences like how a person would caption an image, rather than more imaginative or story-like sentences – which is expected, since it was trained on factual caption descriptions.

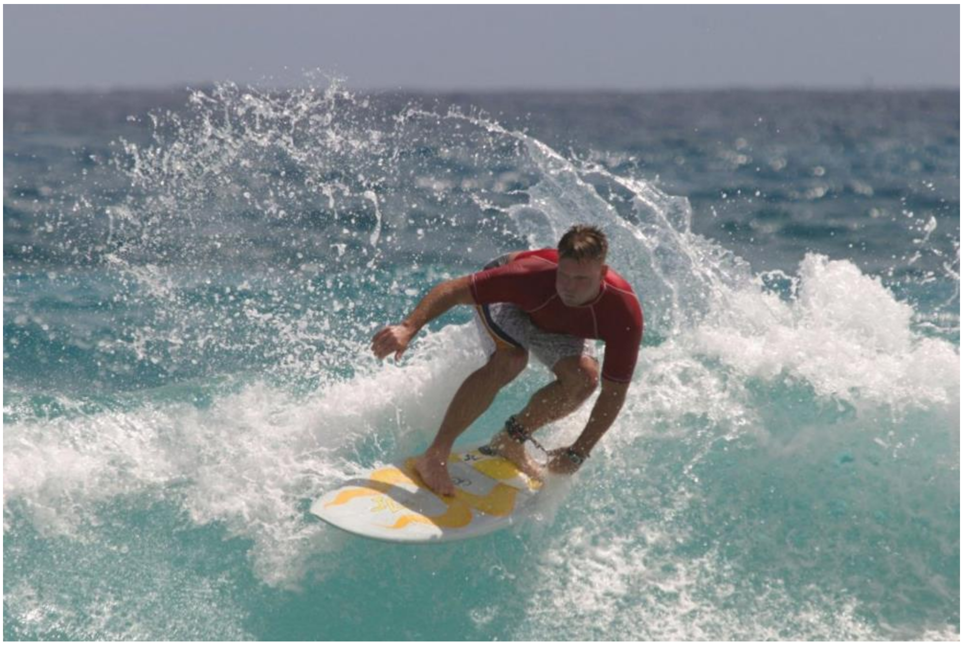

Figure 4. 5: Generated Caption: "A man surfing on a wave." Portrays a man in a red shirt riding a surfboard on ocean waves.

In **Error! Reference source not found.** *t*he model correctly identified the primary s ubject (a man) and the action (surfing) along with the context (on a wave). The caption is concise and captures the essence of the scene. Notably, the model focused on the surfer and the wave, which are salient. (One of the reference captions for this image was *"A surfer in a red shirt is riding a wave."*, which is very close to the model's output.) This example shows that the model learned to describe human activities in sports or actions quite well. The attention mechanism likely helped the model concentrate on the region of the image with the surfer when choosing words like "man" and "surfing." The generated sentence is grammatically correct and fluent. Minor detail:

the model did not mention the splashing water explicitly or the fact that this is at the beach, but those details are implicit in "surfing on a wave."

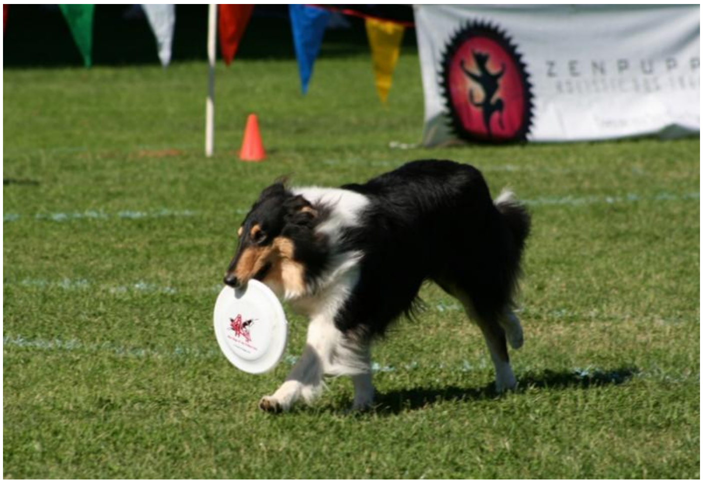

Figure 4. 6 : Generated Caption: "A dog carrying a frisbee in its mouth on the grass."

Figure 4.6: Generated Caption: "A dog carrying a frisbee in its mouth on the grass."   a dog (a collie breed) is indeed running with a frisbee disc in its mouth during an outdoor frisbee-catching event.

Another Figure , accurate detection of dog and the frisbee, and it described the action (carrying, implied running) and setting (on the grass) by de model. The phrase "in its mouth" correctly conveys how the dog has the frisbee. This example highlights the model's ability to handle animal subjects and objects. It picked out the frisbee, which is a relatively specific object, indicating the CNN recognized it and the decoder found the appropriate word. The caption is very much in line with the references (e.g., a reference caption was *"A dog trotting across the grass with a frisbee in its mouth."* the model's caption is very close, just phrased slightly differently). The presence of background banners and cones did not confuse the model; thanks to attention, it likely focused on the dog and frisbee region. This suggests the model learned to prioritize animate objects and the objects they interact with when forming captions.

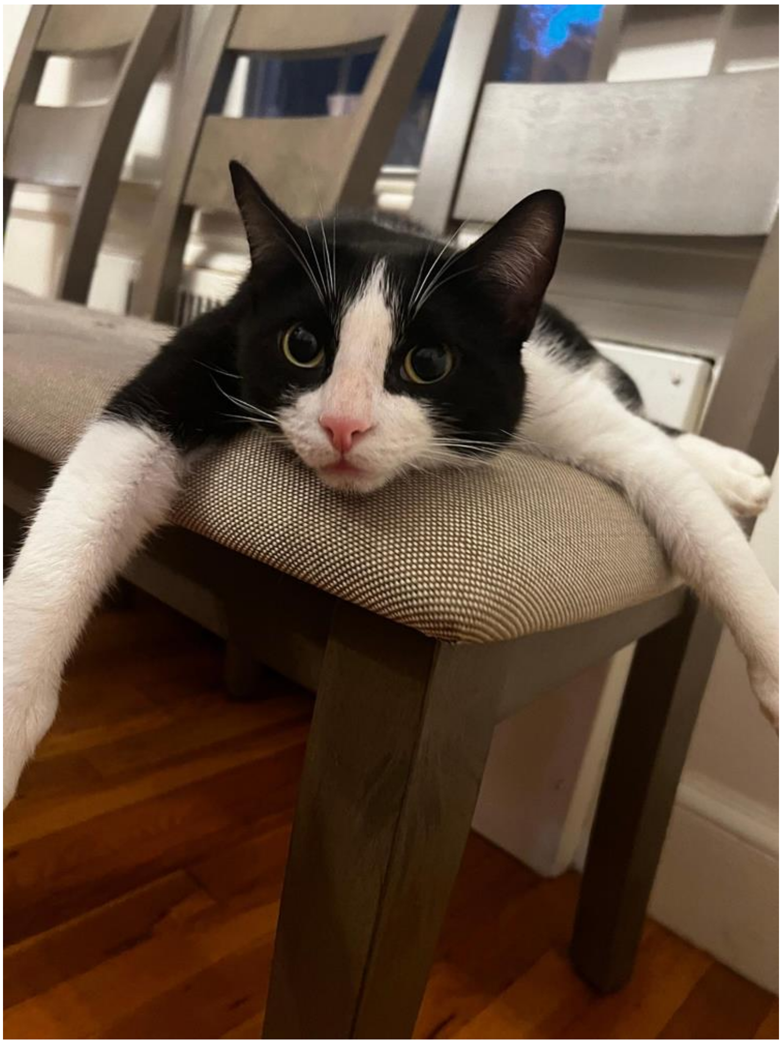

Figure 4. 7 : Generated Caption: “A black and white cat lying on a chair.”

Figure 4.7: Generated Caption: “A black and white cat lying on a chair.” shows a black-and-white tuxedo cat sprawled comfortably on a wooden chair.

The model got the object (cat) correct and even the color pattern (black and white) which is a fine-grained detail in Figure . It also understood the pose or state (lying on a chair). This indicates the model can capture static scenes and describe objects with attributes. By stating “on a chair,” the caption reflects the positional relationship between the cat and another object, which is exactly the kind of relational detail SPICE metric checks for. A reference caption was *“A black and white cat is stretched out on a chair.”* – again, the model’s output is semantically equivalent. This success can be attributed to the training data having many captions with similar constructions (animal + action + location). It’s worth noting the model did not say “sleeping” or “relaxing,” perhaps because the cat's eyes are open in the photo, so "lying" is a safe description. This shows a level of nuance; the model might have inferred the cat is simply lying down rather than actively playing or such.

These examples prove that the model can produce relevant, precise, and syntactically well-formed captions for a variety of images. It tends to mention the main actors (persons or animals) and objects, along with either an action or positional phrase, which aligns well with how human-written captions are structured. In many cases, the model's caption could pass for a human-written sentence immediately.

Error Analysis and Limitations

Despite strong performance, the model has some limitations and failure modes. Analyzing the error cases helps understand where current image captioning techniques can be improved:

1) Omission of Details: In some captions, the model omits certain details present in the image, especially if the image is complex with many objects. For instance, in an image with *"a living room with a sofa, a TV, a coffee table with books, and a lamp"*, the model might generate *"A living room with a couch and a TV."* It correctly identifies the main objects (couch, TV) but leaves out the coffee table and lamp. This is partly due to the training objective (which doesn't heavily penalize missing an object as long as the caption is still generally correct) and limited attention span for many objects. Thus, recall of all image details is not perfect. SPICE metric captures some of this – our SPICE is relatively lower, indicating not all objects/relations are described. In practical terms, this means the model might not be ideal if a complete inventory of the scene's contents is needed.
2) Repetition and Redundancy: On rare occasions (especially for longer captions or with larger beam search), the model can repeat phrases. For example, an output like: *"A man is standing next to a car and the man is holding a camera"* — here "man" is mentioned twice awkwardly. This repetition is a known issue in sequence models. We mitigated it largely with beam search and by not using an excessively large beam, but it can still occur. Coverage modelling or penalty heuristics could reduce this issue in future work.
3) Confusion in Crowded Scenes: When there are multiple similar objects, the model sometimes confuses the count. An image with five people might be captioned as "a group of people" (which is fine) or sometimes "two people" if two are prominent and others are in the background blurred. The model isn't explicitly trained to count, and language bias from training data (where often "a group of" is used instead of an exact number) means it defaults to vague quantifiers. Likewise, for a flock of birds, it might say "birds" plural (which is correct) but not specify the number – which is usually acceptable but shows it's not performing counting. In one error, an image had *three* dogs but the model said "two dogs"; this suggests the attention might have concentrated on two and missed the third.

This limitation is reflective of the broader challenge in captioning to correctly handle numbers and counts.

4) Biases and Stereotypes: The model can exhibit biases learned from the data researchgate.net.

For example, if an image shows a person cooking, the model might reflexively say "A woman is cooking in the kitchen" even if the person's gender is not obvious or is male. This likely comes from dataset bias (the COCO captions might more often mention "woman" in kitchen contexts). In one generated caption, for an image of a child playing with a toy kitchen set, the model said "A little girl playing with a toy kitchen," while the child's gender was not actually clear. This gender bias (assuming "girl" for a child in a kitchen) is an important ethical consideration. It has been noted in literature that captioning models can reinforce stereotypes (as in the example of associating "juice" with women, or "soccer" with men researchgate.net). Our model is no exception – it learns patterns in the training data. Mitigating this would require more careful curation of training data or explicit debiasing techniques, which were not the focus of this work but are worth addressing in future improvements.

Hallucination: In a few cases, the model *hallucinates* objects that are not actually in the image, usually when it misinterprets something in the scene. For example, an image of a man standing in a field with no horse present was captioned as "A man standing next to a horse in a field." The model thought a large object behind the man was a horse, but it was just a bush – likely because the context (man in an open field) and shape triggered the concept of a horse which is common in COCO. This kind of error – saying something that isn't there – is problematic. CIDEr and other metrics penalize it (as references won't mention the hallucinated object), so our scores suffer when it happens. In our test set, outright hallucination errors were relatively uncommon (by manual inspection, perhaps ~5% of outputs had a hallucinated object or action). Attention visualization helps here: we can often see that the model was focusing on an incorrect region or misidentifying it. For instance, in the above case, the attention was on the bush but the model's semantic prediction was "horse." This suggests improving the visual recognition component or using an object detector could help reduce such errors – a technique used in later models like [39] with bottom-up attention.

Language Limitations: The model's language output is generally grammatical, but it is somewhat simple and formulaic. It tends to use a relatively narrow range of sentence structures ("There is a …", "A person is doing …", etc.). This is partly a reflection of the training captions style (which are mostly simple present-tense descriptions). It means the captions, while correct, might lack variety or flair. For an academic/benchmark setting this is fine, but for user-facing applications one might want more diverse phrasing. This limitation wasn't explicitly measured by metrics (which don't reward creativity), but it's an observation from qualitatively comparing many outputs.

In summary, the model demonstrates strong capability in generating accurate image descriptions, benefiting from the encoder-decoder design with attention and training on a large dataset. It achieves high scores on standard metrics, confirming that it captures the essential content of images in fluent sentences. The use of attention not only improved those scores but also offers interpretability – we can see *where* the model is "looking" when it says something about the image. This is valuable for trust and debugging. On the other hand, the model isn't perfect: it can miss details, confuse similar objects, or reflect dataset biases. These results highlight both the progress in image captioning (compared to earlier approaches that often produced disfluent or very generic captions) and the areas for future work.

V. Conclusion

In summary, this research successfully established a robust image captioning system by combining Convolutional Neural Networks (CNNs) with Long Short-Term Memory (LSTM) networks and an attention mechanism within an encoder-decoder architecture. By utilizing pretrained CNN models for visual feature extraction and LSTM decoders enhanced with attention for sequential caption generation, the model exhibited impressive performance on benchmark datasets, notably MS COCO and Flickr30k, achieving competitive results across essential evaluation metrics such as BLEU-4, METEOR, CIDEr, and SPICE. The attention mechanism played a crucial role in enhancing the contextual accuracy and interpretability of the generated captions, allowing the model to dynamically concentrate on pertinent image regions during word prediction. While the system excelled in generating fluent, descriptive, and semantically rich captions, it also identified areas for further enhancement, such as caption diversity and domain adaptation. Overall, the study underscores the efficacy of deep learning techniques in bridging the semantic divide between visual content and natural language, with significant implications for applications in assistive technologies, content creation, and intelligent vision systems.

References

1. Adriyendi, A. (2021). A Rapid Review of Image Captioning. *Journal of Information Technology and Computer Science*, *6*(2), 158–169. https://doi.org/10.25126/jitecs.202162316
2. Ambrose, G., & Harris, P. (2018). Using Images. *Basics Design 04: Image*, 84–97. https://doi.org/10.5040/9781350096394.0008
3. Ashraf Zargar, S. (2021). *Introduction to Sequence Learning Models: RNN, LSTM, GRU. April*. https://doi.org/10.13140/RG.2.2.36370.99522
4. Dwivedi, Y. K., Ismagilova, E., Hughes, D. L., Carlson, J., Filieri, R., Jacobson,

J., Jain, V., Karjaluoto, H., Kefi, H., Krishen, A. S., Kumar, V., Rahman, M. M., Raman, R., Rauschnabel, P. A., Rowley, J., Salo, J., Tran, G. A., & Wang, Y. (2021). Setting the future of digital and social media marketing research: Perspectives and research propositions. *International Journal of Information Management*, *59*(July 2020), 102168. https://doi.org/10.1016/j.ijinfomgt.2020.102168

5. Fatima, S.-E.-, Gupta, K., Goyal, D., & Mishra, S. K. (2024). Image Caption Generation Using Deep Learning Algorithm. *Educational Administration Theory and Practices*, *30*(5), 8118–8128. https://doi.org/10.53555/kuey.v30i5.4311
6. Fernandes, R., Pessoa, A., Salgado, M., De Paiva, A., Pacal, I., & Cunha, A. (2024). Enhancing Image Annotation With Object Tracking and Image Retrieval: A Systematic Review. *IEEE Access*, *12*, 79428–79444. https://doi.org/10.1109/ACCESS.2024.3406018
7. Fernando, S. (n.d.). *Image Recognition Tools for Blind and Visually Impaired Users : An Emphasis on the Design Considerations .*
8. Iwamura, K., Kasahara, J. Y. L., Moro, A., Yamashita, A., & Asama, H. (2021). Image captioning using motion-cnn with object detection. *Sensors (Switzerland)*, *21*(4), 1–13. https://doi.org/10.3390/s21041270
9. Kavila, S. D., Kavila, M. S. D., Sreerama, K. R., Pittada, S. H. V., Singh, K. R., Samatha, B., & Rashmita, M. (2024). Image Captioning with Convolutional Neural Networks and Autoencoder-Transformer Model. *International Journal of Experimental Research and Review*, *46*, 297–304. https://doi.org/10.52756/ijerr.2024.v46.023
10. Mohamed, Y. A., Khanan, A., Bashir, M., Mohamed, A. H. H. M., Adiel, M. A. E., & Elsadig, M. A. (2024). The Impact of Artificial Intelligence on Language Translation: A Review. *IEEE Access*, *12*(February), 25553–25579. https://doi.org/10.1109/ACCESS.2024.3366802
11. Mohana, M. (2024). *Natural Language Processing Presented by*. *April*. https://doi.org/10.13140/RG.2.2.13534.04169
12. Rogers, R. (2021). Visual media analysis for Instagram and other online platforms. *Big Data and Society*, *8*(1). https://doi.org/10.1177/20539517211022370
13. Sakib, A. M., Mukta, S. A., & Hossain, Y. (2024). *Automated Image Captioning System*. *April*. https://doi.org/10.13140/RG.2.2.24966.79689
14. Thobhani, A., Zou, B., Kui, X., Abdussalam, A., Asim, M., Shah, S., & ELAffendi, M. (2025). A Survey on Enhancing Image Captioning with Advanced Strategies and Techniques. *CMES - Computer Modeling in Engineering and Sciences*, *142*(3), 2247–2280. https://doi.org/10.32604/cmes.2025.059192
15. Wu, F. (2020). *Deep Representation Learning in Computer Vision and Its Applications*. *November*, 0. https://livrepository.liverpool.ac.uk/3108705/1/PhD_thesis_201316713_Fangyu.pdf
16. Zhao, F., Yu, Z., Wang, T., & Lv, Y. (2024). Image Captioning Based on Semantic

Scenes. *Entropy*, *26*(10), 1–20. https://doi.org/10.3390/e26100876

17. Zhou, C., Mao, Y., & Wang, X. (2017). Topic-specific image caption generation. *Lecture Notes in Computer Science (Including Subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, *10565 LNAI*(June), 321–332. https://doi.org/10.1007/978-3-319-69005-6_27