

JaSPICE: Automatic Evaluation Metric Using Predicate-Argument Structures for Image Captioning Models

Yuiga Wada
Keio University
yuiga@keio.jp

Kanta Kaneda
Keio University
k.kaneda@keio.jp

Komei Sugiura
Keio University
komei.sugiura@keio.jp

Abstract

Image captioning studies heavily rely on automatic evaluation metrics such as BLEU and METEOR. However, such n -gram-based metrics have been shown to correlate poorly with human evaluation, leading to the proposal of alternative metrics such as SPICE for English; however, no equivalent metrics have been established for other languages. Therefore, in this study, we propose an automatic evaluation metric called JaSPICE, which evaluates Japanese captions based on scene graphs. The proposed method generates a scene graph from dependencies and the predicate-argument structure, and extends the graph using synonyms. We conducted experiments employing 10 image captioning models trained on STAIR Captions and PFN-PIC and constructed the *Shichimi* dataset, which contains 103,170 human evaluations. The results showed that our metric outperformed the baseline metrics for the correlation coefficient with the human evaluation.

1 Introduction

Image captioning has been extensively studied and applied to various applications in society, such as generating fetching instructions for robots, assisting blind people, and answering questions from images (Magassouba et al., 2019; Ogura et al., 2020; Kambara et al., 2021; Gurari et al., 2020; White et al., 2021; Fisch et al., 2020). In this field, it is important that the quality of the generated captions is evaluated appropriately. However, researchers have reported that automatic evaluation metrics based on n -grams do not correlate well with human evaluation (Anderson et al., 2016). Alternative metrics that do not rely on n -grams have been proposed for English (e.g., SPICE (Anderson et al., 2016)); however, they are not fully applicable to all languages. Therefore, developing an automatic evaluation metric that correlates well with human evaluation for image captioning models in languages other than English would be beneficial.

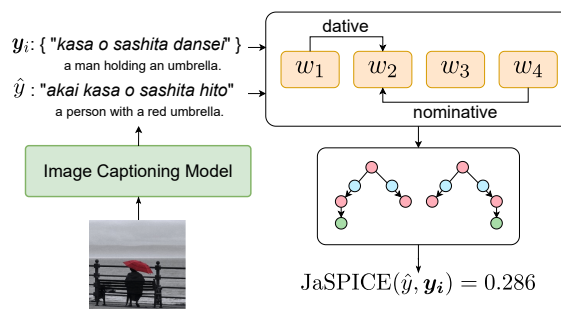


Figure 1: Overview of JaSPICE¹. Given a candidate caption and reference captions, our method parses the scene graph from the PAS and dependencies, and then computes a score that represents the similarity between the candidate and the references by matching both graphs.

SPICE is a standard metric for image captioning in English and evaluates captions based on scene graphs. SPICE uses Universal Dependency (UD) (de Marneffe et al., 2014) to generate scene graphs; however, UD can only extract basic dependencies and cannot handle complex relationships. In the case of Japanese, the phrase “A no B” (Kurohashi et al., 1999), which is composed of the nouns A and B has multiple semantic relations, which makes the semantic analysis of such phrases a challenging problem. For example, in the noun phrase “*kinpatsu no dansei*” (“a blond man”), “blond” (A) is an attribute of “a man” (B) and “a man” (B) is an object, whereas the noun phrase “*dansei no kuruma*” (“man’s car”) represents the relation of a “car” (B) being owned by a “man” (A), and dependency parsing using UD cannot accurately extract the relationship between A and B. Given that UD cannot handle complex relationships and is therefore not suitable for constructing scene graphs, directly applying SPICE to the evaluation of Japanese captions poses challenges. Furthermore, problem settings exist that are difficult to evaluate using SPICE simply computed from an English translation (e.g., TextCaps (Sidorov et al., 2020)).

¹Project page: <https://yuiga.dev/jaspice/en>

To address these issues, we propose JaSPICE, which is an automatic evaluation metric for image captioning models in Japanese. JaSPICE is computed from scene graphs generated from dependencies and the predicate-argument structure (PAS) and can therefore take complex relationships into account.

Fig. 1 illustrates our JaSPICE approach, where the main idea is that we first parse the scene graph from the PAS and dependencies, and then computes a score that represents the similarity between the candidate caption and the reference captions by matching both graphs. For example, given the candidate caption “*akai kasa o sashita hito*” (“a person with a red umbrella”) and the reference caption “*kasa o sashita dansei*” (“a man with an umbrella”), our method parses the scene graph and computes a score by matching both graphs.

Our method differs from existing methods because it generates scene graphs based on dependencies and the PAS and uses synonym sets for the evaluation so that it can evaluate image captioning models in Japanese. It is expected that appropriate scene graphs can be generated by reflecting dependencies and PAS in scene graphs. It is also expected that the use of synonym sets will improve the correlation of metrics with human evaluation because it considers the matching of synonyms that do not match on the surface.

The main contributions are as follows:

- We propose JaSPICE, which is an automatic evaluation metric for image captioning models in Japanese.
- Unlike SPICE which uses UD, JaSPICE generates scene graphs based on dependencies and the PAS.
- We introduce a graph extension using synonym relationships to take synonyms into account in the evaluation.
- We constructed the *Shichimi* dataset, which contains a total of 103,170 human evaluations collected from 500 evaluators.

2 Related Work

2.1 Image Captioning and Its Applications

Many studies have been conducted in the field of image captioning (Xu et al., 2015; Herdade et al., 2019; Cornia et al., 2020; Luo et al., 2021; Ng et al., 2021; Li et al., 2022). For instance, (Stefanini et al., 2021) is a survey paper that provides

a comprehensive overview of image caption generation, including models, standard datasets, and evaluation metrics. Specifically, various automatic evaluation metrics such as embedding-based metrics (Kusner et al., 2015) and learning-based metrics (Zhang et al., 2020) have been comprehensively summarized.

Standard datasets for English image captioning tasks include MS COCO (Lin et al., 2014), Flickr30K (Young et al., 2014) and CC3M (Sharma et al., 2018). Standard datasets for Japanese image captioning tasks include STAIR Captions (Yoshikawa et al., 2017) and YJ Captions (Miyazaki et al., 2016), which are based on MS COCO images.

2.2 Automatic Evaluation Metrics

Standard automatic metrics for image captioning models include BLEU (Papineni et al., 2002), ROUGE (Lin, 2004), METEOR (Banerjee et al., 2005) and CIDEr (Vedantam et al., 2015). Additionally, SPICE (Anderson et al., 2016) is considered as a standard metric for evaluating image captioning models in English.

BLEU and METEOR were first introduced for machine translation. BLEU computes precision using n -grams up to four in length, while METEOR favors the recall of matching unigrams. Additionally, ROUGE considers the longest subsequence of tokens that appears in both the candidate and reference captions, and CIDEr uses the cosine similarity between the TF-IDF weighted n -grams, thereby considering both precision and recall. Unlike these metrics, which are based on n -grams, SPICE evaluates captions using scene graphs.

Scene graph has been widely applied to vision-related tasks such as image retrieval (Johnson et al., 2015; Wang et al., 2020; Schuster et al., 2015), image generation (Johnson et al., 2018), VQA (Beyounes et al., 2019; Li et al., 2019; Shi et al., 2019), and robot planning (Amiri et al., 2022) because of their powerful representation of semantic features of scenes. A scene graph was first proposed in (Johnson et al., 2015) as a data structure for describing objects instances in a scene and relationships between objects. In (Johnson et al., 2015), the authors proposed a method for image retrieval using scene graphs; however, a major shortcoming of their method is that the user needs to enter a query in the form of a scene graph. Therefore, in (Schuster et al., 2015), the authors proposed Stan-

ford Scene Graph Parser, which can parse natural language into scene graphs automatically. (Schuster et al., 2015) is one of the early methods for the construction and application of scene graphs.

SPICE parses captions into scene graphs using Stanford Scene Graph Parser and then computes the F_1 score based on scene graphs. Our method differs from SPICE in that our method introduces a novel scene graph parser based on the PAS and dependencies, graph extensions using synonym relationships so that it can evaluate Japanese captions.

3 Problem Statement

In this study, we focus on the automatic evaluation of image captioning models in Japanese. The terminology used in this study is defined as follows:

- **Predicate-argument structure (PAS):** a structure representing the relation between predicates and their arguments in a sentence.
- **Scene graph:** a graph that represents semantic relations between objects in an image. The details are explained in Section 4.1.

Given a candidate caption \hat{y}_i and a set of reference captions $\{y_{i,j}\}_{j=1}^N$, automatic image captioning evaluation metrics compute a score that captures the similarity between \hat{y}_i and $\{y_{i,j}\}_{j=1}^N$. Note that N denotes the number of reference captions. We evaluate the proposed metric using its correlation coefficient (Pearson/Spearman/Kendall’s correlation coefficient) with human evaluation. This is because automatic evaluation metrics for image captioning models should correlate highly with human evaluation (Anderson et al., 2016).

In this study, we assume that we deal with the automatic evaluation of Japanese image captions. However, some of the discussion in this study can be applied to other languages.

4 Proposed Method

For the evaluation of image captioning models, semantic structure is expected to be more effective than n -gram because, unlike machine translation, image captioning requires grounding based on the scene and relationships between objects in the image. Therefore, utilizing the scene graph, which abstracts the lexical and syntactic aspects of natural language, can be beneficial for the evaluation of image captioning models.

In this study, we propose JaSPICE, which is an automatic evaluation metric for image captioning models in Japanese. JaSPICE is an extension of

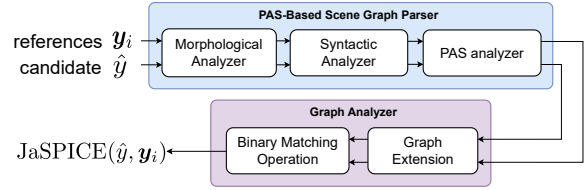


Figure 2: Process diagram of the proposed method. Our method consists of two main modules: PAS-SGP and GA. (i) PAS-SGP generates scene graphs from captions using the PAS and dependencies. (ii) GA performs a graph extension using synonym relationships and then computes the F_1 score by matching tuples extracted from the candidate and the reference scene graphs. JaSPICE is easily interpretable because it outputs the score in the range of $[0, 1]$.

SPICE (Anderson et al., 2016) and can evaluate image captioning models in Japanese based on scene graphs. Although the proposed metric is an extension of SPICE, it also takes into account factors not handled by SPICE, that is, subject completion and the addition of synonymous nodes. Therefore, we believe that the novelty of the proposed metric can be applied to other automatic evaluation metrics.

The main differences between the proposed metric and SPICE are as follows:

- Unlike SPICE, JaSPICE generates a scene graph based on dependencies and the PAS.
- JaSPICE performs heuristic zero anaphora resolution and graph extension using synonyms.

Fig. 2 shows the process diagram of our method. The proposed method consists of two main modules: PAS-Based Scene Graph Parser (PAS-SGP) and Graph Analyzer (GA).

4.1 Scene Graph

The scene graph for a caption y is represented by

$$G(y) = \mathcal{G} \langle O(y), E(y), K(y) \rangle,$$

where $O(y)$, $E(y)$, and $K(y)$ denote the set of objects in y , the set of relations between objects, and the set of objects with attributes, respectively. Given that C , R , and A denote the whole sets of objects, relations, and attributes, respectively, then we can write $O(y) \subseteq C$, $E(y) \subseteq O(y) \times R \times O(y)$, $K(y) \subseteq O(y) \times A$.

Fig. 3 shows an example of an image and scene graph. Fig. 3 (b) shows a scene graph obtained from the description “*hitodōri no sukunaku natta dōro de, aoi zubon o kita otokonoko ga orenji-iro no herumetto o kaburi, sukēto bōdo ni notte iru.*”

(“on a deserted street, a boy in blue pants and an orange helmet rides a skateboard.”) for Fig. 3 (a). The pink, green, and light blue nodes represent objects, attributes, and relationships, respectively, and the arrows represent dependencies.

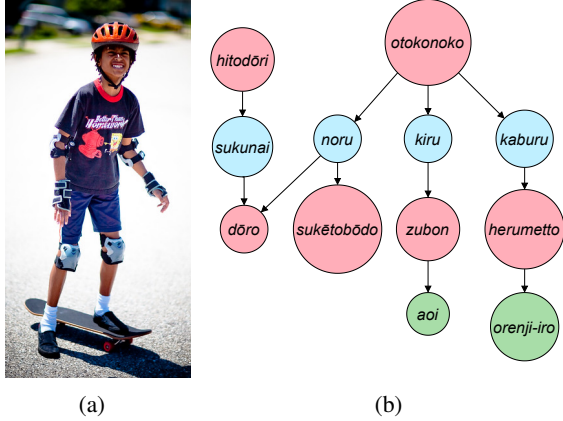


Figure 3: Example of an image and corresponding scene graph. The pink, green, and light blue nodes represent objects, attributes, and relationships, respectively, and the arrows represent dependencies. The caption is “hitodōri no sukunaku natta dōro de, aoi zubon o kita otokonoko ga orenji-iro no herumetto o kaburi, sukēto bōdo ni notte iru.”

4.2 PAS-Based Scene Graph Parser (PAS-SGP)

The input of PAS-SGP is generated caption \hat{y} and the output is scene graph $G(\hat{y})$. First, the morphological analyzer, syntactic analyzer, and predicate-argument structure analyzer² extract the PAS and dependencies from \hat{y} . Next, scene graph $\mathcal{G} \langle O(\hat{y}), E(\hat{y}), K(\hat{y}) \rangle$ is generated from the PAS and the dependencies by a rule-based method based on 10 case markers. Note that the 10 case markers are: *ga*, *wo*, *ni*, *to*, *de*, *kara*, *yor*, *he*, *made* and deep cases (e.g., temporal case) (Kudo et al., 2014). Our parser directly extracts objects, relations, and attributes from the PAS and the dependencies. To parse them, we have defined a total of 13 dependency patterns. These patterns are designed to encapsulate the following constructions and phenomena:

- Subject–object–verb constructions
- Possessive constructions
- Prepositional phrases
- Clausal modifiers of nouns
- Adjectival modifiers
- Postpositional phrases

²In this study, we employed the tools JUMAN++ (Tolmachev et al., 2018) and KNP (Kurohashi et al., 1994).

Furthermore, it is important to consider zero pronouns (Umakoshi et al., 2021) when comparing two sentences. Consider two sentences A and B with the same meaning, but only sentence A contains a zero pronoun. Sentence A contains a relation that includes zero pronouns, which does not match any relation in sentence B. Hence, even though sentences A and B have the same meaning, not all relations match because of the zero pronoun. Therefore, without careful handling, it is not possible to determine a suitable match.

To alleviate this issue, the proposed method performs heuristic zero anaphora resolution. Algorithm 1 shows the node completion algorithm of zero pronouns (ϕ represents a zero pronoun).

Algorithm 1 Node completion for zero pronouns

Input: $r \in R$
for $o \in \text{get_objects}(r)$ **do**
 if $\text{Rel} \langle \phi, r, o \rangle$ is found **then** \triangleright indegree is 0
 for $\text{Rel} \langle o', r', o \rangle \in \text{find_rel}(o)$ **do**
 $\text{Rel} \langle \phi, r, o \rangle \leftarrow \text{Rel} \langle o', r, o \rangle$
 end for
 end if
end for

4.3 Graph Analyzer (GA)

The inputs of GA are $\{G(y_{i,j})\}_{j=1}^N$ and $G(\hat{y})$, where $\{G(y_{i,j})\}_{j=1}^N$ is a set of scene graphs obtained from $\{y_{i,j}\}_{j=1}^N$. First, GA expands $\{G(y_{i,j})\}_{j=1}^N$ and $G(\hat{y})$ by introducing synonym nodes as follows: Suppose that objects o_1 and o_2 are connected by relation r . Given that $S(x)$ denotes the set of synonyms of x , our method generates new relations $\text{Rel} \langle o'_1, r', o'_2 \rangle$, where $o'_1 \in S(o_1)$, $o'_2 \in S(o_2)$, and $r' \in S(r)$. In other words, it adds new nodes $o \in S(o_1) \cup S(o_2)$ and n new edges to the scene graph, where n denotes $(|S(o_1)| + |S(o_2)|) \times |S(r)|$. Note that we use the Japanese WordNet (Bond et al., 2009) to obtain the set of synonyms. We name this process *graph extension*.

Next, GA merges scene graphs $\{G(y_{i,j})\}_{j=1}^N$ into a single graph. Specifically, GA transforms $\mathcal{G} \langle O(y_{i,j}), E(y_{i,j}), K(y_{i,j}) \rangle$ into:

$$G(\mathbf{y}_i) \triangleq \mathcal{G} \langle \{O(y_{i,j})\}_{j=1}^N, \{E(y_{i,j})\}_{j=1}^N, \{K(y_{i,j})\}_{j=1}^N \rangle,$$

where \mathbf{y}_i denotes $\{y_{i,j}\}_{j=1}^N$. To evaluate matching between both scene graphs in the range of $[0, 1]$,

GA computes the F_1 score from $G'(\hat{y})$ and $G(\mathbf{y}_i)$. The F_1 score is appropriate because it can take into account the difference in size between $G'(\hat{y})$ and $G(\mathbf{y}_i)$. Precision P , recall R , and JaSPICE are defined as follows:

$$P(\hat{y}, \mathbf{y}_i) = \frac{|T(G'(\hat{y})) \otimes T(G(\mathbf{y}_i))|}{|T(G'(\hat{y}))|},$$

$$R(\hat{y}, \mathbf{y}_i) = \frac{|T(G'(\hat{y})) \otimes T(G(\mathbf{y}_i))|}{|T(G(\mathbf{y}_i))|},$$

$$\text{JaSPICE}(\hat{y}, \mathbf{y}_i) = \frac{2 \cdot P(\hat{y}, \mathbf{y}_i) \cdot R(\hat{y}, \mathbf{y}_i)}{P(\hat{y}, \mathbf{y}_i) + R(\hat{y}, \mathbf{y}_i)}.$$

Note that we define $T(G(x))$ as:

$$T(G(x)) \triangleq O(x) \cup E(x) \cup K(x),$$

and \otimes denotes a function that returns matching tuples in two scene graphs.

5 Experiments

5.1 Setup

We conducted experiments to compare JaSPICE with existing automatic evaluation metrics. In the experiments, we calculated the correlation coefficients between automatic evaluation metrics and human evaluation. For the evaluation, we used outputs from the image captioning models, $\{y_i\}$ and $\{y_{\text{rand}}\}$, obtained from STAIR Captions (Yoshikawa et al., 2017) and PFN-PIC (Hatori et al., 2018), which consisted of 21,227 and 1,920 captions, respectively. Note that $\{y_i\}$ was randomly selected from $\{y_{i,j}\}_{j=1}^M$, and y_{rand} was randomly selected from all of $\{y_{i,j} | i = 1, \dots, N, j = 1, \dots, M\}$, where M is the number of captions included per image. We used a crowdsourcing service to collect human evaluations from 500 evaluators (The details are explained in Section 5.5). For a given image, the human evaluators rated the appropriateness of its caption on a five-point scale. To evaluate the proposed metric, we calculated the correlation coefficient (Pearson/Spearman/Kendall’s correlation coefficient) between $\{s_J^{(i)}\}_{i=1}^N$ and $\{s_H^{(i)}\}_{i=1}^N$, where $s_J^{(i)}$ and $s_H^{(i)}$ denote the JaSPICE for the i -th caption and the human evaluation for the i -th caption, respectively.

Although there were problems with translation quality and speed, it was technically possible to compute SPICE by translating \hat{y} and $\{y_{i,j} | i = 1, \dots, N, j = 1, \dots, M\}$ into English. Thus, we conducted a comparison experiment between the

Table 1: Correlation coefficients between each automatic evaluation metric and the human evaluation for STAIR Captions.

Metric	Pearson	Spearman	Kendall
BLEU	0.296	0.343	0.260
ROUGE	0.366	0.340	0.258
METEOR	0.345	0.366	0.279
CIDEr	0.312	0.355	0.269
JaSPICE	0.501	0.529	0.413
r_{human}	0.759	0.750	0.669

Table 2: Comparison between JaSPICE and SPICE in terms of correlation with human evaluation for STAIR Captions.

Metric	Pearson	Spearman	Kendall
SPICE _{service}	0.488	0.515	0.402
SPICE _{trm}	0.491	0.516	0.403
JaSPICE	0.501	0.529	0.413

proposed metric and SPICE obtained in this manner. In the experiments, we calculated the correlation coefficient between the human evaluation and SPICE obtained from the English translation. To avoid quality issues specific to a single machine translation, we performed the English translations using multiple approaches. Specifically, we used a vanilla Transformer trained on JParaCrawl (Morishita et al., 2020) and a proprietary machine translation system³.

In this study, we used caption-level correlation $f(\{s_J^{(i)}\}_{i=1}^N, \{s_H^{(i)}\}_{i=1}^N)$ for the evaluation. In (Anderson et al., 2016), caption-level correlation $f(\{s_S^{(i)}\}_{i=1}^N, \{s_H^{(i)}\}_{i=1}^N)$ and system-level correlation $f(\{\bar{s}_S^{(j)}\}_{j=1}^J, \{\bar{s}_H^{(j)}\}_{j=1}^J)$ were used to evaluate the automatic evaluation metric, where f , $s_S^{(i)}$, and J denote the correlation coefficient function, SPICE for the i -th caption and the number of models, respectively. However, because J is generally very small, it is not appropriate to use system-level correlation $f(\{\bar{s}_S^{(j)}\}_{j=1}^J, \{\bar{s}_H^{(j)}\}_{j=1}^J)$ for the evaluation. In fact, in (Kilickaya et al., 2017), the authors also used only the correlation coefficient per caption for the evaluation.

5.2 Corpora and Models

In this study, we used STAIR Captions and PFN-PIC as corpora. STAIR Captions is a large-scale Japanese image-caption corpus, and PFN-PIC is a corpus for a robotic system, which contains object manipulation instructions in English and Japanese.

³We used DeepL as a proprietary machine translation tool.

We adopted these corpora because STAIR Captions is a standard Japanese image caption corpus based on MS-COCO images, and PFN-PIC is a standard dataset that comprises images and a set of instructions in Japanese for a robotic system.

To evaluate the proposed metric on STAIR Captions, we used a set of 10 standard models, including SAT(Xu et al., 2015), ORT(Herdade et al., 2019), \mathcal{M}^2 -Transformer(Cornia et al., 2020), DLCT(Luo et al., 2021), ER-SAN(Li et al., 2022), ClipCap_{mlp}(Mokady et al., 2021), ClipCap_{trm}, and Transformer_{L=3,6,12}(Vaswani et al., 2017). We trained these models on STAIR Captions from scratch. Additionally, to evaluate the proposed metric on PFN-PIC, we used a set of 3 standard models, including CRT(Kambara et al., 2021), ORT, and SAT. The details are explained in Appendix A.

5.3 Experimental Results: STAIR Captions

To validate the proposed metric, we experimentally compared it with the baseline metrics using their correlation with human evaluation.

Table 1 shows the quantitative results for the proposed metric and baseline metrics on STAIR Captions. Note that r_{human} is explained in Section 5.5. For the baseline metrics, we used BLEU(Papineni et al., 2002), ROUGE(Lin, 2004), METEOR(Banerjee et al., 2005) and CIDEr(Vedantam et al., 2015), which are standard automatic evaluation metrics for image captioning.

Table 1 shows that the Pearson, Spearman, and Kendall correlation coefficients between JaSPICE and the human evaluation were 0.501, 0.529 and 0.413, respectively, which indicates that JaSPICE outperformed all the baseline metrics.

Table 2 shows a comparison between JaSPICE and SPICE in terms of correlation with human evaluation. Note that $\text{SPICE}_{\text{trm}}$ and $\text{SPICE}_{\text{service}}$ denote SPICE calculated from English translations by Transformer trained on JParaCrawl and a proprietary machine translation system, respectively. Table 2 indicates that the Pearson, Spearman, and Kendall correlation coefficients between JaSPICE and the human evaluation were 0.501, 0.529 and 0.413, respectively. Thus, JaSPICE outperformed $\text{SPICE}_{\text{trm}}$ by 0.010, 0.013, and 0.010 points for each correlation coefficient, respectively. Similarly, JaSPICE outperformed $\text{SPICE}_{\text{service}}$ by 0.013, 0.014, and 0.011 points.

Fig. 4 show successful examples of the proposed metric for STAIR Captions. Fig. 4 (a) illustrates an input image and its corresponding scene

Table 3: Correlation coefficients between each evaluation metric and human evaluation for PFN-PIC.

Metric	Pearson	Spearman	Kendall
BLEU	0.484	0.466	0.352
ROUGE	0.500	0.474	0.365
METEOR	0.423	0.457	0.352
CIDEr	0.416	0.462	0.353
JaSPICE	0.572	0.587	0.452

graph for \hat{y}_k “*megane o kaketa josei ga aoi denwa o sōsa shite iru*” (“a woman wearing glasses is operating a blue cell phone”). For this sample, $y_{i,1}$ was “*josei ga aoi sumātofon o katate ni motte iru*” (“woman holding blue smartphone in one hand”). Regarding this sample, $\text{JaSPICE}(\hat{y}, \mathbf{y}_i)$ and $s_H^{(i)}$ were 0.588 and 5, respectively. In the STAIR Captions test set, 33.6% of the total samples were rated as $s_H^{(i)} = 5$, whereas the top 33.6% score in $\{\text{JaSPICE}(\hat{y}, \mathbf{y}_k)\}_{k=1}^N$ was observed to be $\tau_S = 0.207$. This sample satisfies $\text{JaSPICE}(\hat{y}, \mathbf{y}_i) > \tau_S$, suggesting that our metric generated an appropriate score for this sample.

Similarly, Fig. 4 (b) shows an input image and scene graph for \hat{y}_j “*akai kasa o sashita hito ga benchi ni suwatte iru*” (“a person with a red umbrella is sitting on a bench”). For Fig. 4 (b), $y_{j,1}$ was “*akai kasa o sashite suwatte umi o mite iru*” (“sitting with a red umbrella, looking out to sea.”), and regarding this sample, $\text{JaSPICE}(\hat{y}, \mathbf{y}_j)$ and $s_H^{(j)}$ were 0.632($> \tau_S$) and 5, respectively. These results indicate that the proposed metric generated appropriate scores for STAIR Captions.

5.4 Experimental Results: PFN-PIC

Table 3 shows the quantitative results for the proposed and baseline metrics for PFN-PIC. Table 3 indicates that the Pearson, Spearman, and Kendall correlation coefficients between JaSPICE and the human evaluation were 0.572, 0.587, and 0.452, respectively, which indicates that JaSPICE outperformed all the baseline metrics.

Table 4 shows the correlation coefficients between JaSPICE and the human evaluation for the PFN-PIC dataset. The results indicate that JaSPICE also outperformed both $\text{SPICE}_{\text{trm}}$ and $\text{SPICE}_{\text{service}}$ on PFN-PIC.

Fig. 5 shows successful examples of the proposed metric for PFN-PIC. Note that the green and red boxes in the figure represent the target object and destination, respectively. Fig. 5 (a) illustrates an input image and its corresponding scene

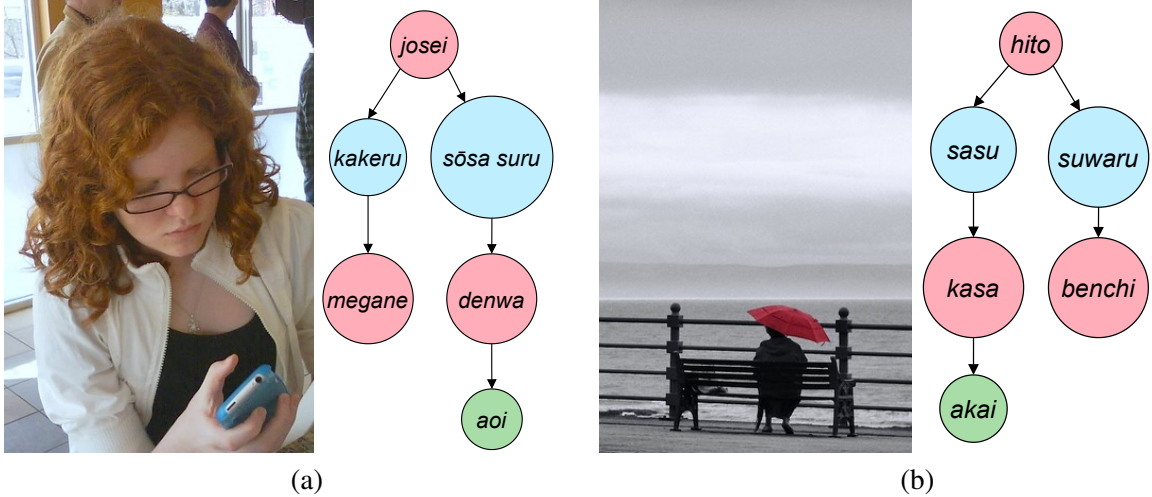


Figure 4: Image and scene graph for successful cases for STAIR Captions. (a) \hat{y}_i : “megane o kaketa josei ga aoi denwa o sōsa shite iru” (“a woman wearing glasses is operating a blue cell phone”), $s_H^{(i)} = 5$, $\text{JaSPICE}(\hat{y}, \mathbf{y}_i) = 0.526 > \tau_S$; and (b) \hat{y}_j : “akai kasa o sashita hito ga benchi ni suwatte iru” (“a person with a red umbrella is sitting on a bench”), $s_H^{(j)} = 5$, $\text{JaSPICE}(\hat{y}, \mathbf{y}_j) = 0.632 > \tau_S$.

Table 4: Comparison of JaSPICE and SPICE in terms of correlation with human evaluation for PFN-PIC.

Metric	Pearson	Spearman	Kendall
SPICE _{service}	0.416	0.418	0.316
SPICE _{trm}	0.427	0.420	0.317
JaSPICE	0.572	0.587	0.452

Table 5: Results of the ablation study (P: Pearson, S: Spearman, K: Kendall, M : the number of samples for which $\text{JaSPICE}(\hat{y}, \mathbf{y}_i) = 0$).

Metric	Parser	Graph Extension	P	S	K	M
(i)	UD		0.398	0.390	0.309	1465
(ii)	UD	✓	0.399	0.390	0.309	1430
(iii)	JaSGP		0.493	0.524	0.410	1417
(iv)	JaSGP	✓	0.501	0.529	0.413	1346

graph for \hat{y}_i “migishita no hako no naka no kōra no kan o, hidariue no hako ni ugokashite kudasai” (“move the can of Coke in the box in the bottom right to the box in the top left”). Regarding Fig. 5 (a), $y_{i,1}$ was “kōra no kan o, hidariue no kēsu ni ugokashite chōdai” (“move the can of Coke to the case in the top left-hand corner”). For this sample, $\text{JaSPICE}(\hat{y}, \mathbf{y}_i)$ and $s_H^{(i)}$ were 0.870 and 5, respectively. In the PFN-PIC test set, 41.2% of the total samples were rated as $s_H^{(i)} = 5$, whereas the top 41.2% score in $\{\text{JaSPICE}(\hat{y}, \mathbf{y}_k)\}_{k=1}^N$ was observed to be $\tau_P = 0.276$. This sample satisfies $\text{JaSPICE}(\hat{y}, \mathbf{y}_i) > \tau_P$, suggesting that our metric generated an appropriate score for this sample.

Similarly, Fig. 5 (b) shows an input image and a scene graph for \hat{y}_j “mizuiro no kappu o, migiue no

hako ni ugokashite kudasai” (“move the blue cup to the box in the top right-hand corner.”). Regarding Fig. 5 (b), $y_{j,1}$ was “hidarishita no hako no naka ni aru mizuiro no kappu o, migiue no hako ni ugokashite kudasai” (“move the blue cup from the bottom left box to the top right box”). For this sample, $\text{JaSPICE}(\hat{y}, \mathbf{y}_j)$ and $s_H^{(j)}$ were 0.385 ($> \tau_P$) and 5, respectively. These results indicate that the proposed metric also generated appropriate scores for PFN-PIC.

5.5 Experimental Results: *Shichimi*

Although the above experiment was compared to baseline metrics, it is also important to compare metrics with r_{human} , the correlation coefficient within human evaluations. Hence, to calculate r_{human} , we constructed the *Shichimi* (Subject Human evaluations of Comprehensive Image captioning Model’s Inferences) dataset containing a total of 103,170 human evaluations collected from 500 evaluators. The *Shichimi* dataset, which includes images, captions, and human evaluations on a five-point scale, is a versatile resource that can be efficiently utilized to develop regression-based metrics such as COMET (Rei et al., 2020).

We found r_{human} to be 0.759 on the *Shichimi* dataset. The reason for r_{human} being less than 1.0 is the variability among human evaluations within the same sample. Here, we define r_{human} as $\mathbb{E}[\mathcal{R}(Y_i, Y_j)]$, where Y_i and \mathcal{R} denote the human evaluation vector by the i -th user and the correlation coefficient function, respectively. r_{human} is considered to be a virtual upper bound on the

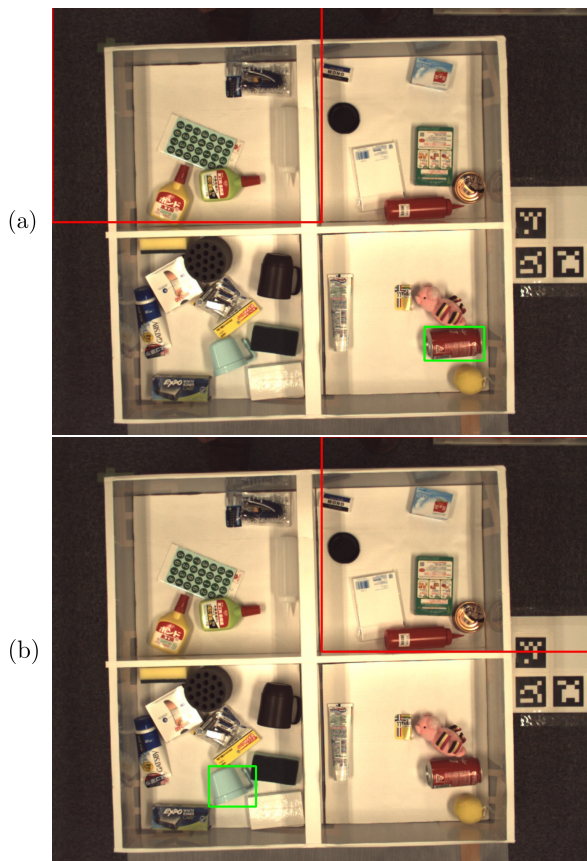


Figure 5: Image for successful cases for PFN-PIC. (The green and red boxes in the figure represent the target object and destination, respectively.) (a) \hat{y}_i : “move the can of Coke in the box in the bottom right to the box in the top left”, $s_H^{(i)} = 5$, $\text{JaSPICE}(\hat{y}, \mathbf{y}_i) = 0.870 > \tau_P$; and (b) \hat{y}_j : “move the blue cup to the box in the top right-hand corner”, $s_H^{(j)} = 5$, $\text{JaSPICE}(\hat{y}, \mathbf{y}_j) = 0.385 > \tau_P$. Scene graphs for these samples are shown in the Appendix C.

performance of the automatic evaluation metrics. Among the baseline metrics, the correlation coefficient of ROUGE, which performed best, was 0.366. This was a difference of 0.393 from r_{human} , indicating that the use of baseline metrics for the evaluation of image captioning could be problematic. Meanwhile, the difference between the correlation coefficient in JaSPICE and r_{human} was 0.258. Although this shows an improvement over the baseline metrics, there remains scope for further enhancement (Error analysis and discussion can be found in Appendix E).

5.6 Ablation Studies

We defined two conditions for ablation studies. Table 5 shows the results of the ablation study. For each condition, we examined not only the correlation coefficient but also the number of samples M for which $\text{JaSPICE}(\hat{y}, \mathbf{y}_i) = 0$. This is be-

cause JaSPICE might produce a zero output when no matched pairs are found during the comparison between pairs in $T(G'(\hat{y}))$ and $T(G(\mathbf{y}_i))$.

Scene Graph Parser Ablation We replaced PAS-SGP with a scene graph parser based on UD (UD parser) to investigate the performance of PAS-SGP. In comparison with Metric (iv), under Metric (ii), the values of the Pearson, Spearman, and Kendall correlation coefficients were 0.102, 0.139, and 0.104 points lower, respectively. Furthermore, there were 119 fewer samples for M . This indicates that the introduction of the PAS-SGP contributed the most to performance.

Graph Extension Ablation We investigated the influence on performance when the graph extension was removed. A comparison between Metric (i) and (iv), in addition to (iii) and (iv), suggests that the introduction of graph extensions also contributed to the performance improvement.

6 Conclusions

In this study, we proposed JaSPICE, which is an automatic evaluation metric for image captioning models in Japanese. The following contributions of this study can be emphasized:

- We proposed JaSPICE, which is an automatic evaluation metric for image captioning models in Japanese.
- Unlike SPICE, we proposed a rule-based scene graph parser PAS-SGP using dependencies and PAS.
- We introduced graph extension using synonyms to take synonyms into account in the evaluation.
- We constructed the *Shichimi* dataset, which contains a total of 103,170 human evaluations collected from 500 evaluators.
- Our method outperformed SPICE calculated from English translations and the baseline metrics on the correlation coefficient with the human evaluation.

In future studies, we will extend our method by taking into account hypernyms and hyponyms.

ACKNOWLEDGMENT

This work was partially supported by JSPS KAKENHI Grant Number 23H03478, JST CREST, and NEDO.

References

- Saeid Amiri, Kishan Chandan, et al. 2022. Reasoning With Scene Graphs for Robot Planning Under Partial Observability. *IEEE RAL*, 7(2):5560–5567.
- Peter Anderson, Basura Fernando, et al. 2016. SPICE: Semantic Propositional Image Caption Evaluation. In *ECCV*, pages 382–398.
- Peter Anderson, Xiaodong He, Chris Buehler, Damien Teney, Mark Johnson, Stephen Gould, and Lei Zhang. 2018. Bottom-Up and Top-Down Attention for Image Captioning and Visual Question Answering. In *CVPR*, pages 6077–6086.
- Satanjeev Banerjee and Alon Lavie. 2005. METEOR: An Automatic Metric for MT Evaluation with Improved Correlation with Human Judgments. In *EEvaluation@ACL*, pages 65–72.
- Hedi Ben-younes, Remi Cadene, Nicolas Thome, and Matthieu Cord. 2019. BLOCK: Bilinear Superdiagonal Fusion for Visual Question Answering and Visual Relationship Detection. In *AAAI*, pages 8102–8109.
- Francis Bond, Hitoshi Isahara, Sanae Fujita, et al. 2009. Enhancing the Japanese WordNet. In *Workshop on Asian Language Resources*, pages 1–8.
- Marcella Cornia, Matteo Stefanini, Lorenzo Baraldi, et al. 2020. Meshed-Memory Transformer for Image Captioning. In *CVPR*, pages 10578–10587.
- MarieCatherine de Marneffe, Timothy Dozat, Natalia Silveira, Katri Haverinen, Filip Ginter, Joakim Nivre, and Christopher D. Manning. 2014. Universal Stanford dependencies: A cross-linguistic typology. In *LREC*, pages 4585–4592.
- Adam Fisch, Kenton Lee, Ming-Wei Chang, Jonathan Clark, and Regina Barzilay. 2020. CapWAP: Image Captioning with a Purpose. In *EMNLP*, pages 8755–8768.
- Danna Gurari, Yinan Zhao, Meng Zhang, and Nilavra Bhattacharya. 2020. Captioning Images Taken by People Who Are Blind. In *ECCV*, pages 417–434.
- Jun Hatori, Yuta Kikuchi, Sosuke Kobayashi, Kuniyuki Takahashi, Yuta Tsuboi, Yuya Unno, et al. 2018. Interactively Picking Real-World Objects with Unconstrained Spoken Language Instructions. In *ICRA*, pages 3774–3781.
- Simao Herdade, Armin Kappeler, et al. 2019. Image Captioning: Transforming Objects into Words. In *NeurIPS*, volume 32, pages 11137–11147.
- Justin Johnson, Agrim Gupta, and Li Fei-Fei. 2018. Image Generation From Scene Graphs. In *CVPR*, pages 1219–1228.
- Justin Johnson, Ranjay Krishna, Michael Stark, Li-Jia Li, David Shamma, et al. 2015. Image Retrieval Using Scene Graphs. In *CVPR*, pages 3668–3678.
- Motonari Kambara and Komei Sugiura. 2021. Case Relation Transformer: A Crossmodal Language Generation Model for Fetching Instructions. *IEEE RAL*, 6(4):8371–8378.
- Mert Kilickaya, Aykut Erdem, Nazli Ikizler-Cinbis, and Erkut Erdem. 2017. Re-evaluating Automatic Metrics for Image Captioning. In *EACL*, pages 199–209.
- Taku Kudo, Hiroshi Ichikawa, and Hideto Kazawa. 2014. A joint inference of deep case analysis and zero subject generation for Japanese-to-English statistical machine translation. In *ACL Short Papers*, pages 557–562.
- Sadao Kurohashi and Makoto Nagao. 1994. A Syntactic Analysis Method of Long Japanese Sentences Based on the Detection of Conjunctive Structures. *Computational Linguistics*, 20(4):507–534.
- Sadao Kurohashi and Yasuyuki Sakai. 1999. Semantic Analysis of Japanese Noun Phrases - A New Approach to Dictionary-Based Understanding. In *ACL*, pages 481–488.
- Matt Kusner, Yu Sun, Nicholas Kolkin, et al. 2015. From Word Embeddings To Document Distances. In *PMLR*, volume 37, pages 957–966.
- Jingyu Li et al. 2022. ER-SAN: Enhanced-Adaptive Relation Self-Attention Network for Image Captioning. In *IJCAI*, pages 1081–1087.
- Linjie Li, Zhe Gan, Yu Cheng, and Jingjing Liu. 2019. Relation-Aware Graph Attention Network for Visual Question Answering. In *ICCV*, pages 10313–10322.
- Chin Lin. 2004. ROUGE: A Package For Automatic Evaluation Of Summaries. In *ACL*, pages 74–81.
- Tsung Lin, Michael Maire, Serge Belongie, Lubomir Bourdev, Ross Girshick, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C. Lawrence Zitnick. 2014. Microsoft COCO: Common Objects in Context. In *ECCV*, pages 740–755.
- Yunpeng Luo, Jiayi Ji, Xiaoshuai Sun, Liujuan Cao, Yongjian Wu, Feiyue Huang, Chia-Wen Lin, and Rongrong Ji. 2021. Dual-Level Collaborative Transformer for Image Captioning. In *AAAI*, volume 35, pages 2286–2293.
- Aly Magassouba et al. 2019. Multimodal Attention Branch Network for Perspective-Free Sentence Generation. In *CoRL*, pages 76–85.
- Takashi Miyazaki and Nobuyuki Shimizu. 2016. Cross-Lingual Image Caption Generation. In *ACL*, pages 1780–1790.
- Ron Mokady, Amir Hertz, and Amit Bermano. 2021. ClipCap: CLIP Prefix for Image Captioning. *arXiv preprint arXiv:2107.06912*.
- Makoto Morishita, Jun Suzuki, and Masaaki Nagata. 2020. JParaCrawl: A Large Scale Web-Based English-Japanese Parallel Corpus. In *LREC*, pages 3603–3609.

- Edwin G. Ng, Bo Pang, Piyush Sharma, and Radu Soricut. 2021. Understanding Guided Image Captioning Performance across Domains. In *CoNLL*, pages 183–193.
- Tadashi Ogura, Aly Magassouba, Komei Sugiura, et al. 2020. Alleviating the Burden of Labeling: Sentence Generation by Attention Branch Encoder-Decoder Network. *IEEE RAL*, 5(4):5945–5952.
- Kishore Papineni, Salim Roukos, Todd Ward, and Wei Zhu. 2002. Bleu: a Method for Automatic Evaluation of Machine Translation. In *ACL*, pages 311–318.
- Ricardo Rei, Craig Stewart, Ana C Farinha, and Alon Lavie. 2020. COMET: A neural framework for MT evaluation. In *EMNLP*, pages 2685–2702.
- Sebastian Schuster, Ranjay Krishna, et al. 2015. Generating Semantically Precise Scene Graphs from Textual Descriptions for Improved Image Retrieval. In *EMNLP 4th Workshop on Vision and Language*, pages 70–80.
- Piyush Sharma, Nan Ding, Sebastian Goodman, and Radu Soricut. 2018. Conceptual captions: A Cleaned, Hypernymed, Image Alt-text Dataset for Automatic Image Captioning. In *ACL*, pages 2556–2565.
- Jiaxin Shi, Hanwang Zhang, and Juanzi Li. 2019. Explainable and explicit visual reasoning over scene graphs. In *CVPR*, pages 8376–8384.
- Oleksii Sidorov, Ronghang Hu, et al. 2020. TextCaps: a Dataset for Image Captioning with Reading Comprehension. In *ECCV*, pages 742–758.
- Matteo Stefanini, Marcella Cornia, Lorenzo Baraldi, Silvia Cascianelli, Giuseppe Fiameni, and Rita Cucchiara. 2021. From Show to Tell: A Survey on Deep Learning-based Image Captioning. *arXiv preprint arXiv:2107.06912*.
- Arseny Tolmachev, Daisuke Kawahara, and Sadao Kurohashi. 2018. Juman++: A morphological analysis toolkit for scriptio continua. In *EMNLP*, pages 54–59.
- Masato Umakoshi, Yugo Murawaki, and Sadao Kurohashi. 2021. Japanese Zero Anaphora Resolution Can Benefit from Parallel Texts Through Neural Transfer Learning. In *EMNLP*, pages 1920–1934.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, et al. 2017. Attention is all you need. In *NeurIPS*, volume 30, pages 5998–6008.
- Ramakrishna Vedantam, Lawrence Zitnick, and Devi Parikh. 2015. CIDEr: Consensus-based Image Description Evaluation. In *CVPR*, pages 4566–4575.
- Sijin Wang, Ruiping Wang, Ziwei Yao, Shiguang Shan, and Xilin Chen. 2020. Cross-modal Scene Graph Matching for Relationship-aware Image-Text Retrieval. In *WACV*, pages 1508–1517.
- Julia White, Gabriel Poesia, Robert Hawkins, Dorsa Sadigh, and Noah Goodman. 2021. Open-domain Clarification Question Generation Without Question Examples. In *EMNLP*, pages 563–570.
- Dustin Wood, P. D. Harms, Graham H. Lowman, and Justin A. DeSimone. 2017. Response Speed and Response Consistency as Mutually Validating Indicators of Data Quality in Online Samples. *Social Psychological and Personality Science*, 8(4):454–464.
- Kelvin Xu, Jimmy Ba, Ryan Kiros, Kyunghyun Cho, Aaron Courville, Ruslan Salakhutdinov, Richard S. Zemel, and Yoshua Bengio. 2015. Show, Attend and Tell: Neural Image Caption Generation with Visual Attention. In *ICML*, pages 2048–2057.
- Yuya Yoshikawa et al. 2017. STAIR Captions: Constructing a Large-Scale Japanese Image Caption Dataset. In *ACL*, pages 417–421.
- Peter Young, Alice Lai, Micah Hodosh, et al. 2014. From Image Descriptions to Visual Denotations: New Similarity Metrics for Semantic Inference over Event Descriptions. *TACL*, 2:67–78.
- Tianyi Zhang, Varsha Kishore, Felix Wu, Kilian Weinberger, and Yoav Artzi. 2020. BERTScore: Evaluating Text Generation with BERT. In *ICLR*.

A Corpora and Systems

The STAIR Captions(Yoshikawa et al., 2017) contains 5 captions for each of 164,062 images, for a total of 820,310 captions. The vocabulary size is 35,642 and the average sentence length is 23.79. The captions were annotated by 2,100 Japanese speakers.

The PFN-PIC(Hatori et al., 2018) is annotated by at least three annotators for each object and divided into training and validation sets. The training set consists of 1,180 images, 25,900 target objects, and 91,590 instructions, and the validation set consists of 20 images, 352 target objects, and 898 instructions.

In the experiments, we divided both STAIR Captions and PFN-PIC into training, validation, and test sets. Note that STAIR Captions included 413,915; 37,269; and 35,594 captions, and PFN-PIC included 81,087; 8,774; and 898 samples, respectively.

To evaluate the proposed metric on STAIR Captions, we used a set of 10 standard models. Table 6 shows the systems used in the experiments. Note that ClipCap_{mlp} and ClipCap_{trm} are variations of ClipCap that incorporate MLP and Transformer as Mapping Networks, respectively, whereas Transformer_L denotes L-layer Transformer models with Bottom-up features(Anderson et al., 2018) as inputs.

Table 6: The system used in the experiments.

System	Citation
SAT	(Xu et al., 2015)
ORT	(Herdade et al., 2019)
Transformer _{L=3}	(Vaswani et al., 2017)
Transformer _{L=6}	(Vaswani et al., 2017)
Transformer _{L=12}	(Vaswani et al., 2017)
M ² -Transformer	(Cornia et al., 2020)
DLCT	(Luo et al., 2021)
ER-SAN	(Li et al., 2022)
ClipCap _{mlp}	(Mokady et al., 2021)
ClipCap _{trm}	(Mokady et al., 2021)
CRT	(Kambara et al., 2021)
Human	—
Random	—

B Applications of image captioning

Numerous studies have been conducted in the field of image captioning(Xu et al., 2015; Herdade et al., 2019; Cornia et al., 2020; Luo et al.,

2021; Li et al., 2022), a crucial area of research that has been further extended and applied in the sphere of robotics(Magassouba et al., 2019; Ogura et al., 2020; Kambara et al., 2021). Multi-ABN(Magassouba et al., 2019) is a model for generating fetching instructions for domestic service robots using multiple images from various viewpoints. ABEN(Ogura et al., 2020) is a model that extends Multi-ABN and introduces linguistic and generative branches to model relationships between subwords, thus achieving subword-level attention. CRT(Kambara et al., 2021) is a model for generating fetching instructions including the spatial referring expressions of target objects and destinations. It introduces Transformer-based encoder-decoder architecture to fuse the visual and geometric features of the objects in images.

C Experimental Results: PFN-PIC

Fig 6 shows the scene graphs for the samples in Fig 5.

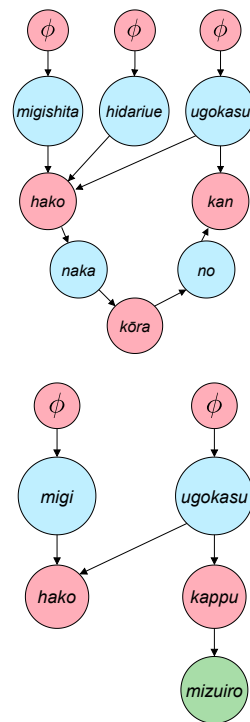


Figure 6: Scene graph for successful cases for PFN-PIC. ϕ represents a zero pronoun. (a) \hat{y}_i : “migishita no hako no naka no kōra no kan o, hidariue no hako ni ugokashite kudasai” (“move the can of Coke in the box in the bottom right to the box in the top left”), $s_H^{(i)} = 5$, JaSPICE(\hat{y}_i, \mathbf{y}_i) = 0.870 > τ_P ; and (b) \hat{y}_j “mizuiro no kappu o, migiue no hako ni ugokashite kudasai” (“move the blue cup to the box in the top right-hand corner”), $s_H^{(j)} = 5$, JaSPICE(\hat{y}_j, \mathbf{y}_j) = 0.385 > τ_P .

D Failure Cases

Fig. 7 shows an unsuccessful example of the proposed metric. Fig. 7 illustrates an input image and its corresponding scene graph for \hat{y}_k “sara ni ryōri ga mora rete iru” (“food is served on a plate”). For Fig. 7, $y_{k,1}$ was “pan ni hamu to kyūri to tomato to chīzu ga hasamatte iru” (“bread with ham, cucumber, tomato and cheese.”). For this sample, JaSPICE was 0 even though $s_H^{(k)}$ was 5. In this case, $y_{k,1}$ used the terms “bread” and “ham” whereas \hat{y} used the hypernym “food”, which resulted in a lower output score because of the mismatch in wording.



Figure 7: Image and scene graph in failed cases for STAIR Captions; \hat{y}_k : “sara ni ryōri ga mora rete iru” (“food is served on a plate”), $s_H^{(k)} = 5$, $\text{JaSPICE}(\hat{y}, y_k) = 0 < \tau_S$.

E Error Analysis and Discussion

We define the failed cases of the proposed metric as a sample that satisfies $\left| \frac{s_H^{(i)}}{\max_i s_H^{(i)}} - \frac{s_J^{(i)}}{\max_i s_J^{(i)}} \right| \geq \theta$. In this study, we set $\theta = 1$ and there were 130 failed samples in the test set.

We investigated 100 out of 130 failed samples. Table 7 categorizes the failure cases. The causes of failure can be divided into five groups:

- (i) Word granularity differences in \hat{y} and y_i : This refers to cases in which y_i used a hyponym for a certain object, relation or attribute in the image, whereas \hat{y} used a hypernym. In the example shown in Fig. 7, the hyponym “bread” was represented by the hypernym “food” in \hat{y} .
- (ii) Difference in focus: This refers to the case in which the focuses of y_i and \hat{y} were different. Both captions were appropriate but focused on different aspects, leading to an inappropriate JaSPICE score.
- (iii) Comparison of sentences containing partially matching morphemes: For example, if \hat{y} was

a sentence containing “tennis racket” and $y_{i,1}$ was a sentence containing “tennis,” then scene graphs had fewer matching pairs, which resulted in an inappropriate JaSPICE.

- (iv) Erroneous evaluation: This refers to cases in which there was a discrepancy between $S_H^{(i)}$ and the quality of \hat{y}_i .
- (v) Others: This category includes other errors.

Table 7 highlights the main bottleneck of the proposed method: the discrepancy in word granularity between \hat{y} and y_i . Therefore, we consider that the bottleneck can be reduced by the introduction of a model that takes into account the relation between hypernyms and hyponyms.

Table 7: Categorization of failed samples.

Error	#Samples
(i)	46
(ii)	20
(iii)	18
(iv)	10
(v) Others	6

F Details of the *Shichimi* Dataset

We removed inappropriate users from the *Shichimi* dataset (e.g. users with extremely short response times (Wood et al., 2017) or those who only responded with the same values).

Table 8 shows the distribution of human evaluations on the *Shichimi* dataset.

Table 8: The distribution on the *Shichimi* dataset.

Score	#Samples
5 (Excellent)	31,809
4 (Good)	21,857
3 (Fair)	22,513
2 (Poor)	12,873
1 (Bad)	14,118