BEiTScore++: Reference-free Image Captioning Evaluation with BEiT-3 and Conformalized Quantile Regression

Gonçalo Gomes[1,2], Bruno Martins[1,2], Chrysoula Zerva[3,4],
[1]Instituto Superior Técnico, University of Lisbon
[2]INESC-ID, [3]Instituto de Telecomunicações
[4]AILS Lab; ECE, National Technical University of Athens
{goncaloecgomes, chrysoula.zerva, bruno.g.martins}@tecnico.ulisboa.pt, czerva@mail.ntua.gr

Abstract

Reliable uncertainty estimation is important for deploying evaluation metrics in real-world settings, where both input ambiguity and domain shifts are common. In this work, we propose an uncertainty-aware extension of the BEiTScore image captioning metric, producing calibrated confidence intervals alongside alignment scores. Our approach models predictive distributions over human judgments, enabling both accurate scoring and principled uncertainty quantification. We evaluate our BEiTScore++ metric across multiple benchmarks with human-annotated ratings, demonstrating strong performance in terms of ranking correlation, calibration, and interval sharpness. BEiTScore++ not only produces accurate predictions, but also captures aleatoric uncertainty arising from ambiguous inputs, and epistemic uncertainty due to limited model knowledge, making it a reliable tool for uncertainty-aware evaluation of image-caption alignment.

1 Introduction

Image captioning evaluation models typically produce a single score that captures the semantic alignment between the image content and a textual description, aiming at scores that are highly correlated with human assessments. However, individual point estimates make it challenging to convey how confident the model is in the produced assessments. This limitation becomes particularly relevant in scenarios where multiple captions receive similar scores while differing substantially in reliability. In these cases, an effective evaluation metric should not only predict caption quality, but also provide an informative estimate of uncertainty.

This work introduces BEiTScore++ as an uncertainty-aware extension to the state-of-the-art BEiTScore image captioning metric (Gomes et al., 2026), which predicts both a quality score and a calibrated confidence interval. We formulate the evaluation problem as quantile regression and train the model using a centered interval objective that encourages both coverage and sharpness. To ensure coherent intervals, the quantiles are constructed monotonically from a base quantile prediction, avoiding quantile crossing by design. Moreover, to improve coverage reliability, we further apply post-hoc conformal calibration. Figure 1 provides a general illustration of the proposed procedure.

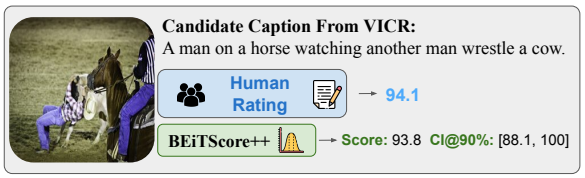


Figure 1: Qualitative example for the use of BEiTScore++ in an instance from the VICR (Narins et al., 2023) benchmark. The candidate caption is scored as highly relevant by humans, as well as by BEiTScore++. Moreover, BEiTScore++ also outputs a confidence interval, covering with high probability the score assigned by humans.

BEiTScore++ is evaluated on several benchmark datasets with human numeric ratings, allowing us to assess how well the predicted uncertainty correlates with prediction error and human judgments. In particular, we show that correlation with human judgments remains at the level of state-of-the-art evaluation metrics, and that the predicted intervals are well calibrated, sufficiently sharp, and informative for downstream analysis.

Beyond predictive performance, we further investigate whether the estimated uncertainty reflects meaningful sources of variability. We focus on two key factors: input ambiguity and model knowledge gaps. To study input ambiguity, we introduce controlled image degradations, such as Gaussian blur, and observe that uncertainty increases proportionally with the level of visual degradation. To examine model knowledge gaps, we evaluate the

Findings of the Association for Computational Linguistics: EMNLP 2026, pages 46192–46226
October 24-29, 2026 ©2026 Association for Computational Linguistics

model on out-of-domain artwork datasets and find that uncertainty rises significantly compared to in-domain natural images, reflecting epistemic uncertainty due to domain shift. Overall, BEiTScore++ aims to make image captioning evaluation not only more accurate, but also more trustworthy.[1]

2 Related Work

This section reviews prior work concerning three key areas: image captioning evaluation, uncertainty quantification, and uncertainty calibration.

2.1 Image Captioning Evaluation

CLIPScore (Hessel et al., 2021) was proposed to overcome the limitations of reference-dependent evaluation methods by enabling reference-free scoring through the cosine similarity between CLIP image and text embeddings. However, the CLIP text encoder typically imposes a strict token limit (e.g., 77 tokens), which restricts its use for long-form descriptions. Several extensions and related approaches have been proposed (Sarto et al., 2025; Hu et al., 2023; Narins et al., 2023; Wada et al., 2024; Gomes et al., 2025b), including approaches that adapt CLIP-based methods to longer textual inputs (Chen et al., 2025; Zhang et al., 2024).

Other approaches leverage Large Language Models (LLMs) for caption evaluation (Chan et al., 2023; Lee et al., 2024; Kim et al., 2025). These models benefit from strong zero-shot capabilities, long context windows, and the ability to capture nuanced semantic differences. However, they incur high computational costs and slow inference times, making them impractical for large-scale benchmarking or integration into decoding pipelines.

More recent work revisits learned encoder-based metrics using more expressive cross-encoder architectures. Gomes et al. (2026) proposed BEiTScore, with basis on a cross-encoder initialized from a BEiT-3 visual question answering checkpoint, this way combining strong pretrained representations with computational efficiency. BEiTScore employs cross-attention between the visual and textual modalities, enabling fine-grained sensitivity to visual-linguistic inconsistencies. Its training strategy incorporates a diverse data mixture designed to expose the model to long descriptions and complex visual-linguistic relationships, along with adversarial augmentations that enhance compositional coverage and relational reasoning.

[1]The corresponding source code and the different models are available at https://github.com/gecgomes/Uncertainty_BEiTScore_pp

Empirical results show that BEiTScore maintains robust compositional understanding across both short and long descriptions, while matching or surpassing specialized LLM-based metrics at a fraction of the computational cost, making it a viable and cost-effective alternative to LLM-based caption evaluation approaches.

2.2 Uncertainty Quantification

Applications that rely on machine learning predictions often require an estimate of confidence alongside the prediction itself, in order to mitigate overconfident decisions. The process of estimating this confidence is commonly referred to as uncertainty quantification. From a philosophical perspective, uncertainty can be framed by asking *where does the randomness lie?* This leads to two distinct viewpoints, namely Bayesian approaches which attribute randomness to model parameters, and frequentist approaches which assume randomness originates from the data.

In Bayesian methods, neural network weights are modeled as probability distributions rather than fixed values (Mackay, 1992; Welling and Teh, 2011; Tran et al., 2019). Training focuses on learning these distributions and, during inference, weights are sampled accordingly. While theoretically appealing, these approaches are computationally expensive, particularly for modern models with millions or billions of parameters, leading to the widespread adoption of frequentist approximations.

A common approximation-based approach corresponds to deep ensembles, where multiple models are trained independently and the empirical variance of their predictions is used as a measure of uncertainty (Garmash and Monz, 2016; Kendall and Gal, 2017; Pearce et al., 2020; Zhan et al., 2023). Another widely used technique is Monte Carlo dropout (Gal and Ghahramani, 2016; Glushkova et al., 2021), which applies dropout at test time to perform multiple stochastic forward passes, estimating uncertainty through variations in the outputs. However, both approaches require multiple forward passes at inference time to construct uncertainty estimates, which can become increasingly slow as the number of samples grows.

Beyond variance-based approaches, quantile regression has emerged as a compelling alternative, particularly in settings where the model can be fully trained and controlled. Quantiles offer a flexible and non-parametric representation of uncertainty, are directly interpretable in the target space, and en-

able straightforward construction of prediction intervals. Traditionally, quantile regression is trained using the pinball loss (Koenker and Hallock, 2001). However, recent work (Chung et al., 2021) shows that the trade-off between sharpness and coverage that is induced by the pinball loss is implicit and often suboptimal, proving that, in highly expressive models, this trade-off tends to favor sharpness at the expense of proper calibration. To address this limitation, Chung et al. (2021) propose alternative loss functions that explicitly balance sharpness and coverage calibration during training.

2.3 Uncertainty Calibration

Uncertainty calibration naturally complements uncertainty quantification, as confidence estimates are only meaningful if they reflect the true frequency with which events occur. In the context of prediction intervals, calibration requires that nominal confidence levels match empirical coverage, for example aiming toward a 90% interval that contains the true value approximately 90% of the time.

For quantile-based methods, a key prerequisite for well-calibrated intervals is to avoid *quantile crossing*, where lower quantiles exceed higher ones and the induced prediction sets become incoherent. Narayan et al. (2021) address this by enforcing monotonicity in the quantile level within deep lattice networks, guaranteeing non-crossing quantiles while retaining model flexibility.

Conformal Risk Control (CRC) provides a general, distribution-free framework for calibrating arbitrary user-defined risk functions. Angelopoulos et al. (2022) extend split conformal prediction by choosing a scalar parameter that controls the expected value of any bounded loss that is monotone in the size of the prediction set. When the loss is chosen as an indicator of miscoverage, CRC recovers classical conformal prediction with finite-sample coverage guarantees, while more general loss choices enable direct control of application-specific risks, such as the false negative rate or task-specific distortion measures.

Recent work has, for instance, leveraged this framework to improve both granularity and calibration of image captioning evaluation metrics. Gomes et al. (2025a) generate distributions over CLIPScore values via simple perturbation strategies (e.g., input masking), and then apply conformal risk control to calibrate these score distributions for word-level error detection and to construct calibrated CLIPScore intervals. This approach yields formal guarantees on user-specified risk bounds, detects foil words in captions, and improves correlation between uncertainty estimates and residuals, thus enhancing CLIPScore's reliability as an image captioning evaluation metric.

3 Properties of Uncertainty Estimates

Uncertainty estimates should satisfy several desirable properties that make them both statistically valid and practically useful. In this work, we focus on four central criteria: calibration, sharpness, monotonicity, and conditional coverage. We also emphasize interpretability and robustness to distribution shift, since uncertainty is only useful if it can be understood and trusted in deployment.

Calibration means that nominal confidence levels match empirical frequencies: for example, a 90% prediction interval should contain the true target about 90% of the time. In turn, sharpness refers to the informativeness of the estimate, typically measured by the width of the prediction interval or the dispersion of the predictive distribution. Following standard principles in probabilistic forecasting, one should first aim to achieve valid coverage and then, among calibrated estimators, maximize sharpness (Gneiting and Raftery, 2007; Chung et al., 2021). In other words, narrow intervals are desirable only when they remain reliable.

For quantile-based uncertainty estimates, the predicted quantiles should be monotone in the quantile level. This prevents quantile crossing, where a lower quantile exceeds a higher one, rendering the resulting interval incoherent. Enforcing non-crossing is important not only for mathematical consistency, but also for producing interpretable confidence intervals that can be reliably used downstream (Narayan et al., 2021).

A calibrated uncertainty estimate may satisfy marginal coverage, meaning that the desired coverage holds on average over the full data distribution. However, marginal validity alone can hide systematic failures on particular regions of the input space or on minority subgroups (Angelopoulos and Bates, 2021; Zerva and Martins, 2024). Conditional coverage is therefore a stronger and more desirable property, requiring reliable coverage within subpopulations or for individual types of inputs. In practice, it is often the difference between an uncertainty estimate that is globally valid and one that is actually dependable for decision-making in heterogeneous settings. However, conditional coverage is something that, in most cases, is very hard if

Algorithm 1 Monotone quantile construction for uncertainty-aware predictions in BEiTScore++

Require: Image I and text description T

$\mathbf{h} \leftarrow \text{BEiT3}(I, T)$

Predict median and bounded deltas:

$q_{\text{mid}} \leftarrow \sigma(f_{\text{mid}}(\mathbf{h}))$

$\boldsymbol{\delta}_{\text{low}} \leftarrow \sigma(f_{\text{low}}(\mathbf{h})) \in [0,1]^{K_{\text{low}}}$

$\boldsymbol{\delta}_{\text{high}} \leftarrow \sigma(f_{\text{high}}(\mathbf{h})) \in [0,1]^{K_{\text{high}}}$

Initialize quantile list: $\mathbf{q} \leftarrow [\, q_{\text{mid}} \,]$

// Build lower tail quantiles by prepending

for $i = 1$ to K_{low} **do**

  $q_{\text{new}} = \delta_{\text{low}}^{(K_{\text{low}}-i+1)} \cdot \mathbf{q}[0]$

  Prepend q_{new} to $\mathbf{q}$

// Build upper tail quantiles by appending

for $j = 1$ to K_{high} **do**

  $q_{\text{new}} = \mathbf{q}[-1] + \delta_{\text{high}}^{(j)} \cdot (1 - \mathbf{q}[-1])$

  Append q_{new} to $\mathbf{q}$

return $\mathbf{y} \leftarrow \mathbf{q}$

not impossible to achieve (Vovk, 2012), and thus a relaxation of this property is often targeted.

Uncertainty estimates should also be easy to interpret by users and practitioners. This is especially important when the output is a confidence interval, quantile set, or calibrated score, since the representation should communicate not only how uncertain the model is but also what the uncertainty means in the scale of the task. Interpretable uncertainty is more likely to support human review, debugging, and downstream risk-aware decisions.

Finally, a useful uncertainty estimate should be sensitive to different types of inputs. In particular, uncertainty should typically increase when the model encounters inputs that are highly ambiguous, or that differ from the training distribution, rather than having outputs that remain overconfident.

4 An Uncertainty Aware BEiTScore

We formulate uncertainty estimation for BEiTScore as a quantile regression problem. The proposed model, which we named BEiTScore++, combines a BEiT-3 backbone with a quantile regression head comprising 11 output nodes, which predict five centered prediction intervals (70%, 75%, 80%, 85%, 90%) together with a central value corresponding to the 50th percentile. These intervals provide a nested family of "typical" uncertainty bands around the median. Jointly learning this grid of quantiles encourages the model to capture a smooth, coherent approximation of the full conditional distribution of scores, which acts as a form of distributional regularization compared to training independent models for individual quantiles (Narayan et al., 2021).

Quantile Regression Loss: Rather than relying on the standard pinball loss, we use the Centered Interval Loss proposed by Chung et al. (2021), which explicitly balances coverage calibration and interval sharpness. We denote a $(1-\alpha)$ confidence interval by $\hat{l}_\alpha = \mathbb{Q}_{\frac{\alpha}{2}}$ and $\hat{u}_\alpha = \mathbb{Q}_{1-\frac{\alpha}{2}}$, and define the centered interval loss as:

$$\mathcal{L}_{\text{CI}}^{\alpha} = (\hat{u}_\alpha - \hat{l}_\alpha) + \tfrac{2}{\alpha}(\hat{l}_\alpha - y)\,\mathbb{I}\{y < \hat{l}_\alpha\} + \tfrac{2}{\alpha}(y - \hat{u}_\alpha)\,\mathbb{I}\{y > \hat{u}_\alpha\}. \quad (1)$$

This loss explicitly penalizes the width of the intervals, thus promoting sharpness. Chung et al. (2021) show that the minimum of its expected value is attained at the true conditional quantiles, thereby also encouraging well-calibrated coverage. The considered design reflects the goal of probabilistic forecasting: intervals should be not only narrow but also correctly calibrated. Accordingly, the lower and upper quantiles are trained jointly to improve interval quality, while the median estimate is supervised with an L_1 loss. The full training objective is thus the sum of the centered interval losses across the five intervals, plus the L_1 loss on the median, defined as follows:

$$\mathcal{L}_{\text{total}} = \sum_{\alpha \in \mathcal{A}} \mathcal{L}_{\text{CI}}^{\alpha} + \lambda \left| y - \mathbb{Q}_{0.5} \right|, \quad (2)$$

where $\mathcal{A} = \{0.10, 0.15, 0.20, 0.25, 0.30\}$ corresponds to the $(90\%, 85\%, 80\%, 75\%, 70\%)$ prediction intervals. Figure 2 illustrates the proposed architecture for uncertainty modeling, and the corresponding post-hoc calibration procedure.

Monotone Quantile Construction. The middle quantile in BEiTScore++, i.e., the 50th percentile, is predicted directly by the respective node, and the remaining quantiles are constructed iteratively from the center by predicting positive deltas towards the lower and upper tail. Concretely, the model first estimates the central score, and then progressively expands toward lower and higher quantiles using bounded multiplicative offsets. This design enforces quantile monotonicity by construction, thus preventing quantile crossing and ensuring that the predicted intervals remain ordered. In addition, because all quantile outputs are constrained to lie in $[0, 1]$, they remain compatible with the score range of BEiTScore and preserve interpretability. Algorithm 1 summarizes the monotone quantile construction used in BEiTScore++.

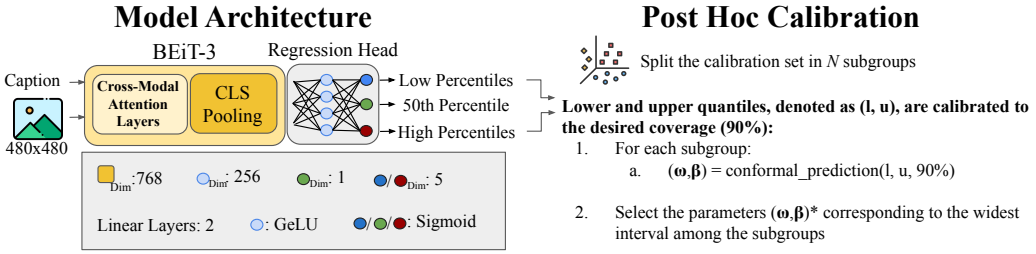


Figure 2: The proposed architecture for uncertainty modeling and the respective post-hoc calibration method.

Quality Score Estimation. The regression head that predict both the median and the other quantile estimates are initialized from the original BEiTScore model, which was trained with a diverse data mixture designed to expose the model to long descriptions and complex visual-linguistic relations. In case of the nodes that estimate the lower and upper deltas, we additionally considered a bias over the weights that pushes the lower quantiles toward values near 0, and the upper quantiles toward values near 1. This scheme which corresponds to subtracting a 2.5 from the bias associated to the original BEiTScore regression head, provides an initial state where the model produces wide intervals with high coverage, from which training can progressively refine calibration and sharpness. This initialization follows the principle that probabilistic forecasts should first achieve calibration and then optimize sharpness, as emphasized by Gneiting and Raftery (2007). Empirically, we found that initializing all quantile heads identically and aiming at the generation of small values (i.e., maximum sharpness), without this specific change in the bias, led to a suboptimal local minimum, producing intervals that were very sharp but poorly calibrated, with weaker correlation between prediction error and estimated uncertainty.

Conformalizing Intervals. We apply a post-hoc calibration step using conformal risk control. We define the calibrated prediction interval as follows:

$$\hat{l} = \max\big(0, (l - m) \cdot \omega + m - \beta\big), \quad (3)$$

$$\hat{u} = \max\big(0, (u - m) \cdot \omega + m + \beta\big), \quad (4)$$

where l, m, and u denote the lower, median, and upper quantile predictions that are being targeted for conformalization, respectively, and where ω and β are calibration parameters optimized within the conformal framework to match a desired coverage level. Our procedure follows a general conformal risk control formulation (Angelopoulos et al., 2022; Gomes et al., 2025a), where the calibration objective is to satisfy a user-specified risk constraint, in our case corresponding to miscoverage.

Approximating Conditional Coverage: As noted by Angelopoulos and Bates (2021), conformal calibration on a held-out calibration set provides formal guarantees only for marginal coverage. While this is valuable, it may hide poor behavior on specific subsets of the data. To obtain a more robust calibration that better approximates conditional coverage, we group training and calibration examples according to image similarity. Specifically, we extracted image embeddings and apply k-means clustering to partition the data into homogeneous subgroups. The number of clusters (i.e., k = 63 in our tests) was selected using a combination of the elbow method and the silhouette score. We then partitioned the training and calibration set according to these clusters.

With the training dataset partitioned into subgroups, we organize mini-batches so that, at each iteration, all samples in a batch are drawn from a single subgroup selected at random. This design is inspired by the group-batching strategy used to induce adversarial group calibration in Chung et al. (2021), where datapoints are grouped based on input features and gradient updates are computed on group-specific batches to improve adversarial group calibration. In our setting, we additionally use the partitioned calibration set to perform conformal calibration independently within each subgroup, obtaining N candidate pairs of calibration parameters (ω, β). We then select the most conservative pair (i.e., the one producing the widest calibrated interval) so that the resulting calibration is valid across all clusters. This strategy is intended to improve robustness across subpopulations by prioritizing coverage preservation, even at the cost of slightly wider intervals.

5 Experimental Evaluation

This section starts by presenting the evaluation methodology, describing implementation details,

datasets, and evaluation metrics. It then describes the experimental results, starting with tests that assessed the correlation to human ratings, and then analyzing the uncertainty estimates.

5.1 Implementation Details

We implement BEiTScore++ using a simple regression head on top of BEiT-3 encoders of size base and large. The regression head is composed of two hidden layers with 256 units each and GeLU activation functions. The encoder weights, together with the output nodes, are initialized from the original BEiTScore model. In turn, the upper/lower quantile nodes are additionally initialized by subtracting 2.5 from the original bias, as described in Section 4. This value was chosen empirically so that the quantiles start with a coverage closer to their respective targets. The complete output layer consists of 11 nodes with sigmoid activations. The model is trained with a batch size of 16 using early stopping with a patience of 15 epochs. All experiments were conducted on a single NVIDIA H200 GPU, although comparable results can be obtained on lesser-memory GPUs by reducing the batch size and using gradient accumulation accordingly.

5.2 Datasets and Evaluation Metrics

Model training uses a mixture of datasets with human-annotated image-caption alignment scores, including VICR (Narins et al., 2023), Polaris (Wada et al., 2024), Nebula (Matsuda et al., 2024), and EvalMuse-40k (Han et al., 2024). These datasets provide continuous human judgments which we normalize to the range [0,1].

To evaluate predictive performance and the quality of uncertainty estimates, we conduct experiments on several benchmark datasets with human ratings. These include VICR (Narins et al., 2023), Polaris (Wada et al., 2024), Nebula (Matsuda et al., 2024), Flickr8K-Expert (Hodosh et al., 2013), Flickr8K-CF (Hodosh et al., 2013), and Composite (Aditya et al., 2015). In addition, for the analysis of different sources of uncertainty, we consider additional datasets spanning in-domain and out-of-domain distributions. In-domain datasets include IIW (Garg et al., 2024), DOCCI (Onoe et al., 2024), Localized Narratives (Pont-Tuset et al., 2020), and MSCOCO (Lin et al., 2014). In turn, out-of-domain datasets focus on artwork images, namely with Docent (Ananthram et al., 2025) and ArtCaps (Lu et al., 2022). Appendix A provides a detailed characterization of all the datasets that were used in the experiments.

We assess model performance using four complementary metrics: coverage, Kendall's τ_c, sharpness, and UPS. Coverage evaluates whether predicted intervals achieve the desired confidence level. Kendall's τ_c measures ordinal agreement between median predictions and human ratings, being widely used in prior work focused on image captioning evaluation metrics. Sharpness quantifies the tightness of the predicted intervals, with narrower intervals indicating more informative uncertainty estimates. Given the predicted lower and upper interval bounds l_i and u_i for each of the N evaluated samples, sharpness is computed as the mean interval width,

$$\text{Sharpness} = \frac{1}{N}\sum_{i=1}^{N}(u_i - l_i)\,. \tag{5}$$

Notice that this measure is only meaningful when the predicted quantiles are monotone. In the presence of quantile crossing, where $l_i > u_i$ for some samples, the width $u_i - l_i$ becomes negative for these samples and no longer corresponds to a valid interval. For these samples sharpness ceases to be a reliable measure of informativeness and can be driven artificially low by crossed predictions.

Finally, the Uncertainty Pearson Score (UPS) measures the correlation between predicted uncertainty and prediction error, capturing whether higher uncertainty aligns with larger deviations from human judgments. Taking the interval width $\sigma_i = u_i - l_i$ as the uncertainty estimate and $e_i = |m_i - y_i|$ as the absolute prediction error, where m_i is the 50th quantile and y_i is the human rating, UPS is defined as the Pearson correlation between the two:

$$\text{UPS} = \rho\,(e, \sigma)\,, \tag{6}$$

where $\rho(\cdot,\cdot)$ denotes the Pearson correlation coefficient, and $e = (e_1, \ldots, e_N)$ and $\sigma = (\sigma_1, \ldots, \sigma_N)$ collect the per-sample errors and uncertainties. Higher UPS values indicate that the model assigns greater uncertainty precisely to the samples on which its predictions are least accurate.

5.3 Assessments Using Human Ratings

Our main goal is to assess whether the estimated uncertainty, measured by the width of the predicted confidence intervals, correlates well with the model's prediction error, defined as the absolute difference between the human rating and the predicted 50th percentile. In parallel, we measure Kendall's

Prior to Calibration				VICR				Polaris				Nebula				Flickr8K-Expert				Composite			
Model	Head	Loss	Clust.	τ	UPS	Sharp	Cov	τ	UPS	Sharp	Cov	τ	UPS	Sharp	Cov	τ	UPS	Sharp	Cov	τ	UPS	Sharp	Cov
Before Calibration																							
MC Dropout-B	-	-	No	78.2	40.2	0.23	68.3	61.5	57.9	0.15	62.4	55.3	39.1	0.25	65.8	59.4	38.5	0.17	64.6	58.9	15.3	0.3	46.3
MC Dropout-L	-	-	No	78.8	42.1	0.18	63.2	63.2	60.8	0.11	68.4	56.7	42.9	0.21	61.0	60.2	48.0	0.12	58.8	59.5	13.7	0.24	42.6
BEiTS++-B	Unc.	PinB.	No	71.7	14.4	0.02	3.5	58.7	22.2	0.01	0.3	52.1	10.8	0.01	13.9	45.1	15.6	0.02	2.9	59.9	12.4	0.04	18.5
BEiTS++-B	Ch.	PinB.	No	78.4	42.9	0.27	60.1	61.6	58.7	0.16	33.9	56.0	47.5	0.31	49.8	59.8	53.1	0.20	34.9	58.8	25.5	0.36	42.2
BEiTS++-B	Ch.	CI	No	78.5	43.9	0.28	61.5	61.4	59.8	0.16	34.2	56.1	48.0	0.31	50.9	59.9	54.5	0.21	36.6	59.0	25.8	0.37	43.8
BEiTS++-B	Ch.	CI	Yes	78.5	43.8	0.28	61.4	61.5	59.7	0.16	34.1	56.1	47.9	0.31	50.6	59.9	54.2	0.21	36.4	59.0	25.9	0.37	43.5
BEiTS++-L	Ch.	CI	Yes	79.0	46.2	0.27	59.1	63.0	57.5	0.15	34.2	57.2	44.2	0.30	50.1	60.7	55.8	0.19	32.8	59.4	24.8	0.34	41.5
After Calibration																							
MC Dropout-B	-	-	No	78.2	43.5	0.34	85.5	61.5	**60.6**	0.22	97.3	55.3	42.1	0.38	84.7	59.4	45.8	0.23	76.8	58.9	20.0	0.45	68.7
MC Dropout-L	-	-	No	78.8	45.6	**0.29**	83.1	**63.2**	60.2	**0.17**	97.0	56.7	43.4	**0.34**	85.2	60.2	50.4	**0.19**	73.4	59.5	21.5	**0.39**	65.0
BEiTS++-B	Unc.	PinB.	No	71.7	11.8	0.49	**94.5**	58.7	18.3	0.48	95.4	52.1	7.9	0.49	83.7	45.1	12.6	0.49	**94.8**	**59.9**	10.0	0.51	**81.0**
BEiTS++-B	Ch.	PinB.	No	78.4	42.8	0.30	86.6	61.6	58.6	**0.17**	96.7	56.0	47.4	**0.34**	79.4	59.8	53.0	0.22	77.6	58.8	25.5	0.40	61.1
BEiTS++-B	Ch.	CI	No	78.5	43.9	0.31	86.9	61.4	59.8	0.18	97.0	56.1	**48.0**	**0.34**	78.9	59.9	54.4	0.23	79.0	59.0	25.9	0.41	59.2
BEiTS++-B	Ch.	CI	Yes	78.5	43.8	0.34	90.5	61.5	59.5	0.21	**97.6**	56.1	47.9	0.38	**85.6**	59.9	54.1	0.26	82.8	59.0	**26.0**	0.45	69.9
BEiTS++-L	Ch.	CI	Yes	**79.0**	**46.1**	0.30	86.8	63.0	57.4	0.18	97.4	**57.2**	44.1	**0.34**	81.2	**60.7**	**55.7**	0.22	77.1	59.4	24.8	**0.39**	63.7

Table 1: Quantitative results across multiple datasets and model configurations. **Bold** indicates the best value and underline indicates the second-best value for each metric within each dataset. For sharpness, lower values are better.

τ_c between the resulting median scores and the human ratings. Ideally, we want BEiTScore++ to match the state-of-the-art correlation with human judgment results of the original BEiTScore metric. We also analyze sharpness, to determine whether improved uncertainty estimation comes at the cost of wider intervals. Together, these metrics allow us to evaluate whether the model achieves a useful trade-off between calibration, sharpness, correlation with human judgments, and the extent to which higher uncertainty is associated with larger deviations from human ratings.

Table 1 presents a quantitative comparison across multiple datasets and model configurations, including a MC dropout baseline over the original BEiTScore, the proposed uncertainty-aware BEiTScore++ variants, and a set of ablations that justify our design choices. Specifically, the table includes: (i) MC dropout as a uncertainty baseline, evaluated with the BEiTScore backbone in both base and large variants; (ii) a variant where the quantile regression head predicts quantiles without the proposed chained mechanism to avoid quantile crossing, trained with the standard pinball loss; and (iii) a variant that does not use our subgroup-based batching strategy for adversarial group calibration, nor the subgroup-wise calibration procedure for relaxed conditional coverage calibration.

We observe consistent improvements when replacing the standard loss and the regression head with our centered-interval loss paired with chained quantile heads, particularly in UPS and sharpness across all benchmarks. After post-hoc calibration, coverage also improves across datasets. Furthermore, comparing models trained with and without clustering and applying subgroup-wise calibration shows incremental gains in the target metrics before calibration, and substantial gains in coverage afterward, especially on the out-of-distribution benchmarks (i.e., the Flickr8K-Expert and Composite datasets). Compared with Monte Carlo dropout, our method achieves better uncertainty quality while remaining considerably more efficient at inference, since it does not rely on repeated forward passes (i.e., 25 passes in the case of our tests). Importantly, the median-score correlation of BEiTScore++ matches that of the original BEiTScore across the evaluated benchmarks, indicating that introducing uncertainty estimation does not degrade the underlying scoring capability. In other words, BEiTScore++ preserves the state-of-the-art human-correlation performance of BEiTScore, while providing calibrated and more informative uncertainty estimates.

A further consequence of omitting the chained mechanism concerns how the intervals can be calibrated. When the quantile regression head predicts each quantile independently, we observe quantile crossing. This prevents us from applying the same two-parameter conformal calibration of Equations 3 and 4. In the presence of crossing the multiplicative factor becomes ill-defined, so we are forced to set ω to zero and calibrate only β, which reduces Equations 3 and 4 to a symmetric band around the median. Beyond the well-documented problems that quantile crossing introduces in quantile regression, restricting the calibration to this single degree of freedom yields poorer calibrated intervals that are also wider, consistent with the sharpness and coverage degradation reported for this ablation in Table 1.

In complement to the results presented in Table 1, Table 2 provides a more detailed evaluation of the correlation between the median pre-

Model	Size (M)	GFLOPs	VICRτ_c	Polarisτ_c	Nebulaτ_c	EX-8Kτ_c	COMτ_c
LLM-based Methods							
FLEUR (Lee et al., 2024)	13350.84	9317.47	74.7	58.4	52.7	53.0	<u>60.8</u>
EXPERT (Kim et al., 2025)	13350.84	9317.47	77.4	60.9	54.9	56.7	**61.8**
VQAScore (CLIP-FlanT5-XXL) (Lin et al., 2024)	11459.59	3742.10	69.8	51.6	49.4	41.0	53.4
Encoder-based Methods							
CLIP-S/ViT-B (Hessel et al., 2021)	151.28	7.28	70.8	52.9	47.8	52.3	51.8
CLIP-S/ViT-L (Hessel et al., 2021)	427.62	84.4	72.6	53.2	48.8	54.7	55.4
LongCLIP-S/ViT-B (Zhang et al., 2024)	149.84	26.24	70.2	54.8	48.9	53.0	51.3
LongCLIP-S/ViT-L (Zhang et al., 2024)	427.75	98.90	69.0	53.9	48.6	53.9	50.5
PAC-S/ViT-B (Sarto et al., 2023)	151.28	4.88	71.4	52.3	47.2	54.3	53.5
PAC-S/ViT-L (Sarto et al., 2023)	427.62	56.26	73.2	53.1	47.9	55.7	55.4
PAC-S++/ViT-B (Sarto et al., 2025)	155.42	4.88	72.7	52.4	46.9	54.5	55.7
PAC-S++/ViT-L (Sarto et al., 2025)	430.58	56.26	75.2	53.6	48.9	57.4	59.4
SPECS (Chen et al., 2025)	151.49	9.20	33.4	39.7	34.4	30.4	24.1
BEiTScore-B/32	222.28	83.85	78.2	61.5	55.3	59.4	58.9
BEiTScore-L/14	674.69	296.73	<u>78.8</u>	<u>63.2</u>	<u>56.7</u>	<u>60.2</u>	59.5
BEiTScore++-B/32	222.28	83.85	78.5	61.5	56.1	59.9	59.0
BEiTScore++-L/14	674.69	296.73	**79.0**	**63.0**	**57.2**	**60.7**	59.4

Table 2: Comparison against state-of-the-art reference-free metrics across multiple human correlation benchmarks. Methods are categorized as encoder-based or LLM-based. Performance is measured using Kendall's τ_c correlation coefficient with human judgments on all benchmarks. The highest score is highlighted using **bold**, and the second best is highlighted using an <u>underline</u>.

dicted scores and human judgments, across a wide range of benchmarks. We compare the proposed uncertainty-aware BEiTScore++ metric against several representative reference-free metrics from the literature, including both encoder-based and LLM-based approaches. This analysis is intended to verify that introducing uncertainty estimation does not compromise the model's ability to produce accurate quality assessments.

Overall, the results confirm that our uncertainty-aware extension retains state-of-the-art performance of BEiTScore in terms of agreement with human judgments. In particular, the large version of the proposed model achieves the best or second-best performance on most benchmarks, outperforming many existing alternatives while providing the additional information about the uncertainty of the model in its predictions.

5.4 Different Sources of Uncertainty

This section investigates whether the uncertainty estimates produced by BEiTScore++ meaningfully reflect different underlying sources of uncertainty. We focus on two complementary factors: input ambiguity, arising from degradation or corruption of the input images, and model knowledge gaps, stemming from exposure to ambiguous and out-of-domain data. Through controlled experiments, we assess whether the uncertainty estimator responds in a predictable and interpretable manner. Results are reported for the large variant, although the base-sized model yields similar conclusions, with results being provided in Appendix D.

5.4.1 Input Ambiguity

To examine whether BEiTScore++ responds appropriately to input ambiguity, we designed an experiment in which images from the Polaris test split were subjected to varying levels of Gaussian blur. The underlying hypothesis is that blur degrades the visual information available to the model, and consequently the model uncertainty should increase proportionally to the level of blur applied.

A key preprocessing step was taken to avoid confounding effects arising from the score distribution. The original Polaris test split is heavily skewed towards low scores, with a large proportion of instances receiving a human rating of zero. For already poor image-caption pairs, the addition of blur could plausibly reinforce the model's confidence in assigning low scores, rather than inducing uncertainty through wider intervals. This behavior would obscure the relationship we aim to measure. To mitigate this, we downsampled instances with a human score of zero, and built a subset of 3624 instances with a more uniform distribution across the full spectrum of scores, aiming for the effect of blur on model uncertainty to be observed clearly across instances of varying quality.

With the balanced subset, we applied progressively stronger levels of blur to each image and measured the corresponding model uncertainty at each level. Figure 3 shows BEiTScore++ interval width distributions under different blur perturbation levels. The results show that our method behaves as expected. The width of the confidence intervals increases with blur intensity set at 10, and this result is statistically significant. An additional interesting behavior emerges at blur level 20, where a reduction in interval width is observed. We attribute this to the images becoming so severely distorted that the model grows confident that the image and

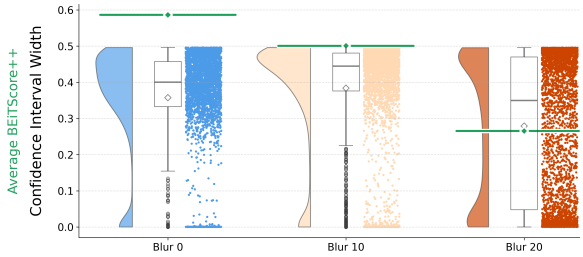


Figure 3: Distribution of 90% confidence interval widths generated by the large variant of BEiTScore++, as Gaussian blur increases on images from the Polaris test split. The horizontal lines indicate the average median score produced by BEiTScore++ at each blur level.

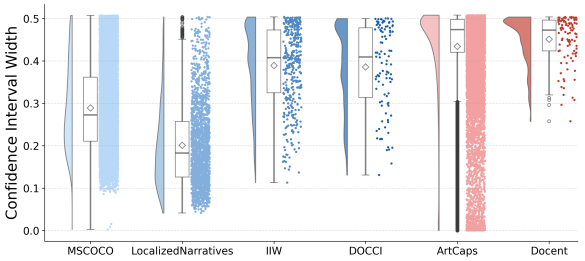


Figure 4: Distribution of 90% confidence interval widths produced by the large variant of the BEiTScore++ metric across in-domain images and out-of-domain artwork datasets, and also across short- and long-caption settings. The color gradient indicates the caption length, with pastel tones for short captions, and more saturated tones for detailed descriptions.

caption are misaligned, effectively starting to collapse uncertainty in favor of a low-score prediction. In addition to interval width distribution, Figure 3 also presents the average BEiTScore++ ratings for each blur level, confirming that the model produces lower scores for increased input degradation.

We additionally evaluated model behavior under a broader set of image distortions, including salt-and-pepper noise, JPEG compression artifacts, random patch occlusions, elastic warping (via spatial smoothing and edge curling), and contrast adjustments (see Appendix C for full details). Across most distortion types, we observed a consistent pattern analogous to Gaussian blur. As the severity of the corruption increased, the model's predictive uncertainty grew proportionally, until we reach a very high level of distortion and the model starts to confidently reject the image-caption alignment. This further supports the hypothesis that the estimator is sensitive to input ambiguity, and that the model's confidence is properly reported.

5.4.2 Model Knowledge Gaps

To study uncertainty arising from model knowledge gaps, we analyze the model's confidence on in-domain datasets with natural images, including long-caption datasets IIW (Garg et al., 2024), DOCCI (Onoe et al., 2024), and Localized Narratives (Pont-Tuset et al., 2020), as well as the short-caption dataset MSCOCO (Lin et al., 2014), and compare them against artwork datasets Docent (Ananthram et al., 2025) and ArtCaps (Lu et al., 2022). These artwork datasets are not only out-of-domain but also more ambiguous: their images are paintings rather than realistic photographs, with stylized shapes, smoother edges, and greater abstraction, while the captions, especially the detailed descriptions in Docent, often involve a higher degree of subjective interpretation.

Since BEiTScore++ supports both detailed and shorter captions, we divide the evaluation into two complementary groups: long captions, in which we compare IIW, DOCCI, and Localized Narratives against Docent; short captions, in which we compare ArtCaps against MSCOCO.

Figure 4 indicates that in both the short- and long-caption settings, the model exhibits higher uncertainty when evaluating artwork datasets. This behavior reflects the fact that artwork imagery can be more challenging to interpret, and also an increased epistemic uncertainty due to a model knowledge gap, as the model was not fine-tuned on artistic imagery. Model uncertainty also increases with the average size of the captions.

6 Conclusions

We introduced an uncertainty-aware extension of BEiTScore that predicts caption quality and generates calibrated confidence intervals. By combining quantile regression with a monotone construction of quantiles, together with a centered interval loss, BEiTScore++ achieves a balance between calibration and sharpness. Post-hoc conformal calibration further improves coverage, particularly when extended with subgroup-based calibration to better approximate conditional coverage.

Empirical results demonstrate that our method maintains strong agreement with human judgments, while producing uncertainty estimates that are well-calibrated. Uncertainty correlates with prediction error and adapts to input ambiguity, indicating that the model captures meaningful signals about its own reliability. These properties are essential for downstream applications, where evaluation metrics inform benchmarking or model selection.

Limitations and Ethical Considerations

While BEiTScore++ improves uncertainty estimation for image-caption evaluation, providing calibrated and sharp confidence intervals, several important limitations remain.

First, our experiments are restricted to English captions, so the behavior of the method in multilingual settings is still unknown. Extending BEiTScore++ to other languages is an important direction for future work, especially given recent evidence that CLIP-based metrics can be adapted for multilingual caption evaluation (Gomes et al., 2025b). It will be important to study whether the proposed uncertainty estimates remain reliable when captions contain translation artifacts, language-specific constructions, or morphological complexity that differs from English.

Second, although our method improves evaluation interpretability, it does not remove the need for human evaluation. Automated metrics can support and complement human judgment, but they should not be treated as a substitute, particularly in applications where erroneous assessments may have downstream consequences. More generally, uncertainty intervals may help users identify low-confidence cases, but they do not guarantee fairness or correctness in all settings.

Third, the proposed framework is evaluated on benchmark datasets that may not fully capture the diversity of real-world image-caption pairs. As with many evaluation methods, its performance may be affected by dataset bias, domain shift, or limited coverage of rare concepts and styles. Future work should investigate how BEiTScore++ behaves under broader distributional shifts, and whether additional calibration strategies are needed to maintain robustness in more challenging settings.

Finally, while BEiTScore++ is designed to be efficient at inference time, the broader research area still depends on substantial computational resources for training and evaluation. In that sense, developing uncertainty-aware metrics that remain accurate, efficient, and broadly applicable across languages and domains remains as a particularly important research direction.

We also note that we used GitHub Copilot during the development of our research work, and we used ChatGPT for minor verifications during the preparation of this manuscript.

Acknowledgements

We thank the anonymous reviewers for their valuable comments and suggestions. This research was supported by the Portuguese Recovery and Resilience Plan through project C645008882-00000055 (i.e., the Center For Responsible AI), and by Fundação para a Ciência e Tecnologia (FCT) through the PhD scholarship with reference 2025.06139.BD, and through the projects with references 2024.07385.IACDC and UIDB/50021/2020 (DOI:10.54499/UIDB/50021/2020). In addition, C. Zerva was partly supported by the framework of the Pharos AI Factory project, funded by the European High-Performance Computing Joint Undertaking (EuroHPC JU) under Grant Agreement No. 101234269 as part of the Horizon Europe and by the Greek Public Investments Program programme, as well as by a Google Research Scholar award.

References

Somak Aditya, Yezhou Yang, Chitta Baral, Cornelia Fermuller, and Yiannis Aloimonos. 2015. From Images to Sentences Through Scene Description Graphs Using Commonsense Reasoning and Knowledge. *arXiv preprint arXiv:1511.03292*.

Amith Ananthram, Elias Stengel-Eskin, Lorena A Bradford, Julia Demarest, Adam Purvis, Keith Krut, Robert Stein, Rina Elster Pantalony, Mohit Bansal, and Kathleen McKeown. 2025. PoSh: Using Scene Graphs To Guide LLMs-as-a-Judge For Detailed Image Descriptions. *arXiv preprint arXiv:2510.19060*.

Anastasios N. Angelopoulos and Stephen Bates. 2021. A Gentle Introduction to Conformal Prediction and Distribution-Free Uncertainty Quantification. *arXiv preprint arXiv:2107.07511*.

Anastasios N. Angelopoulos, Stephen Bates, Adam Fisch, Lihua Lei, and Tal Schuster. 2022. Conformal Risk Control. *arXiv preprint arXiv:2208.02814*.

David Chan, Suzanne Petryk, Joseph Gonzalez, Trevor Darrell, and John Canny. 2023. CLAIR: Evaluating Image Captions with Large Language Models. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing*.

Xiaofu Chen, Israfel Salazar, and Yova Kementchedjhieva. 2025. SPECS: Specificity-Enhanced CLIP-Score for Long Image Caption Evaluation. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing*.

Youngseog Chung, Willie Neiswanger, Ian Char, and Jeff Schneider. 2021. Beyond Pinball Loss: Quantile

Methods for Calibrated Uncertainty Quantification. *In Proceedings of the Annual Meeting on Neural Information Processing Systems*.

Yarin Gal and Zoubin Ghahramani. 2016. Dropout as a Bayesian Approximation: Representing Model Uncertainty in Deep Learning. In *Proceedings of the International Conference on Machine Learning*.

Roopal Garg, Andrea Burns, Burcu Karagol-Ayan, Yonatan Bitton, Ceslee Montgomery, Yasumasa Onoe, Andrew Bunner, Ranjay Krishna, Jason Michael Baldridge, and Radu Soricut. 2024. ImageInWords: Unlocking Hyper-Detailed Image Descriptions. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing*.

Ekaterina Garmash and Christof Monz. 2016. Ensemble Learning for Multi-Source Neural Machine Translation. In *Proceedings of the International Conference on Computational Linguistics*.

Taisiya Glushkova, Chrysoula Zerva, Ricardo Rei, and André F. T. Martins. 2021. Uncertainty-Aware Machine Translation Evaluation. In *Findings of the Conference on Empirical Methods in Natural Language Processing*.

Tilmann Gneiting and Adrian E. Raftery. 2007. Strictly Proper Scoring Rules, Prediction, and Estimation. *Journal of the American Statistical Association*.

Goncalo Emanuel Cavaco Gomes, Bruno Martins, and Chrysoula Zerva. 2025a. A Conformal Risk Control Framework for Granular Word Assessment and Uncertainty Calibration of CLIPScore Quality Estimates. In *Findings of the Association for Computational Linguistics*.

Goncalo Emanuel Cavaco Gomes, Bruno Martins, and Chrysoula Zerva. 2026. BEiTScore: Reference-free Image Captioning Evaluation with an Efficient Cross-Encoder Model. *arXiv preprint arXiv:2605.21728*.

Goncalo Emanuel Cavaco Gomes, Chrysoula Zerva, and Bruno Martins. 2025b. Evaluation of Multilingual Image Captioning: How far can we get with CLIP models? In *Findings of the Association for Computational Linguistics*.

Shuhao Han, Haotian Fan, Jiachen Fu, Liang Li, Tao Li, Junhui Cui, Yunqiu Wang, Yang Tai, Jingwei Sun, Chunle Guo, and 1 others. 2024. Evalmuse-40k: A Reliable and Fine-Grained Benchmark with Comprehensive Human Annotations for Text-to-Image Generation Model Evaluation. *arXiv preprint arXiv:2412.18150*.

Jack Hessel, Ari Holtzman, Maxwell Forbes, Ronan Le Bras, and Yejin Choi. 2021. CLIPscore: A Reference-Free Evaluation Metric for Image Captioning. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing*.

Micah Hodosh, Peter Young, and Julia Hockenmaier. 2013. Framing Image Description as a Ranking Task: Data, Models and Evaluation Metrics. *Journal of Artificial Intelligence Research*.

Anwen Hu, Shizhe Chen, Liang Zhang, and Qin Jin. 2023. InfoMetIC: An Informative Metric for Reference-free Image Caption Evaluation. In *Proceedings of the Annual Meeting of the Association for Computational Linguistics*.

Alex Kendall and Yarin Gal. 2017. What Uncertainties Do We Need in Bayesian Deep Learning for Computer Vision? *In Proceedings of the Annual Meeting on Neural Information Processing Systems*.

Hyunjong Kim, Sangyeop Kim, Jongheon Jeong, Yeongjae Cho, and Sungzoon Cho. 2025. EXPERT: An Explainable Image Captioning Evaluation Metric with Structured Explanations. In *Findings of the Annual Meeting of the Association for Computational Linguistics*.

Roger Koenker and Kevin F. Hallock. 2001. Quantile regression. *Journal of Economic Perspectives*, 15(4).

Yebin Lee, Imseong Park, and Myungjoo Kang. 2024. FLEUR: An Explainable Reference-Free Evaluation Metric for Image Captioning Using a Large Multimodal Model. In *Proceedings of the Annual Meeting of the Association for Computational Linguistics*.

Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. 2014. Microsoft COCO: Common Objects in Context. In *Proceedings of the European Conference on Computer Vision*.

Zhiqiu Lin, Deepak Pathak, Baiqi Li, Jiayao Li, Xide Xia, Graham Neubig, Pengchuan Zhang, and Deva Ramanan. 2024. Evaluating Text-to-Visual Generation with Image-to-Text Generation. In *Proceedings of the European Conference on Computer Vision*.

Yue Lu, Chao Guo, Xingyuan Dai, and Fei-Yue Wang. 2022. ArtCap: A Dataset for Image Captioning of Fine Art Paintings. *IEEE Transactions on Computational Social Systems*, 11(1).

David John Cameron Mackay. 1992. *Bayesian Methods for Adaptive Models*. California Institute of Technology.

Kazuki Matsuda, Yuiga Wada, and Komei Sugiura. 2024. DENEB: A Hallucination-Robust Automatic Evaluation Metric for Image Captioning. In *Proceedings of the Asian Conference on Computer Vision*.

Taman Narayan, Serena Wang, Kevin Canini, and Maya Gupta. 2021. Regularization Strategies for Quantile Regression. *arXiv preprint arXiv:2102.05135*.

Lothar D. Narins, Andrew Scott, Aakash Gautam, Anagha Kulkarni, Mar Castanon, Benjamin Kao, Shasta Ihorn, Yue-Ting Siu, James M. Mason,

Alexander Blum, and 1 others. 2023. Validated Image Caption Rating Dataset. *In Proceedings of the Annual Meeting on in Neural Information Processing Systems*.

Yasumasa Onoe, Sunayana Rane, Zachary Berger, Yonatan Bitton, Jaemin Cho, Roopal Garg, Alexander Ku, Zarana Parekh, Jordi Pont-Tuset, Garrett Tanzer, and 1 others. 2024. DOCCI: Descriptions of Connected and Contrasting Images. In *Proceedings of the European Conference on Computer Vision*.

Tim Pearce, Felix Leibfried, and Alexandra Brintrup. 2020. Uncertainty in Neural Networks: Approximately Bayesian Ensembling. In *Proceedings of the International Conference on Artificial Intelligence and Statistics*.

Jordi Pont-Tuset, Jasper Uijlings, Soravit Changpinyo, Radu Soricut, and Vittorio Ferrari. 2020. Connecting Vision and Language with Localized Narratives. In *Proceedings of the European Conference on Computer Vision*.

Sara Sarto, Manuele Barraco, Marcella Cornia, Lorenzo Baraldi, and Rita Cucchiara. 2023. Positive-Augmented Contrastive Learning for Image and Video Captioning Evaluation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*.

Sara Sarto, Nicholas Moratelli, Marcella Cornia, Lorenzo Baraldi, and Rita Cucchiara. 2025. Positive-Augmented Contrastive Learning for Vision-and-Language Evaluation and Training. *International Journal of Computer Vision*, 133(11).

Dustin Tran, Mike Dusenberry, Mark Van Der Wilk, and Danijar Hafner. 2019. Bayesian Layers: A Module for Neural Network Uncertainty. *In Proceedings of the Annual Meeting on Neural Information Processing Systems*.

Vladimir Vovk. 2012. Conditional Validity of Inductive Conformal Predictors. In *Proceedings of the Asian Conference on Machine Learning*.

Yuiga Wada, Kanta Kaneda, Daichi Saito, and Komei Sugiura. 2024. POLOS: Multimodal Metric Learning from Human Feedback for Image Captioning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*.

Max Welling and Yee W. Teh. 2011. Bayesian Learning via Stochastic Gradient Langevin Dynamics. In *Proceedings of the International Conference on Machine Learning*.

Chrysoula Zerva and André F. T. Martins. 2024. "Conformalizing Machine Translation Evaluation". *Transactions of the Association for Computational Linguistics*, 12:1460–1478.

Runzhe Zhan, Xuebo Liu, Derek F. Wong, Cuilian Zhang, Lidia S. Chao, and Min Zhang. 2023. Test-time Adaptation for Machine Translation Evaluation by Uncertainty Minimization. In *Proceedings of the Annual Meeting of the Association for Computational Linguistics*.

Beichen Zhang, Pan Zhang, Xiaoyi Dong, Yuhang Zang, and Jiaqi Wang. 2024. Long-CLIP: Unlocking the Long-Text Capability of CLIP. In *Proceedings of the European Conference on Computer Vision*.

A Datasets

The following datasets were used in the training, calibration, and evaluation of BEiTScore++, specifically in experiments that focused on assessing the Uncertainty Pearson Score (UPS) and the correlation with human judgments. We separate the datasets into those used for training and calibration, and those used only for evaluation.

Datasets used for training and calibration.

- **VICR** (Narins et al., 2023): The Validated Image Caption Rating (VICR) dataset features human ratings collected through a gamified approach for gathering image-caption pairs, involving a total of 9,990 distinct images. The authors of the dataset demonstrated that it exhibits a superior inter-rater agreement compared to other alternatives (e.g., an improvement of 19% in Fleiss' κ when compared to the agreement for the Flickr8K-Expert dataset), and it features a more balanced distribution across various levels of caption quality. The dataset is organized into 10,175 training pairs, 2,310 validation pairs, and 3,161 test pairs, with an average of 4.4 annotations per instance. In our experiments, we used the test split of the VICR dataset, which includes 3,161 unique image-caption pairs. For calibration, we used the original validation split, which comprises 2,310 unique image-caption pairs.
- **Polaris** (Wada et al., 2024): The Polaris dataset comprises human judgments on image-caption pairs involving 13,691 distinct images, collected from evaluators. The dataset is organized into 26,340 training pairs, 8,738 validation pairs, and 8,726 test pairs, with an average of 3.0 annotations per instance. Polaris includes captions generated by ten standard image captioning models, covering both modern and older architectures to ensure output diversity. In our tests, we used the test split of the Polaris dataset, which includes 8,726 unique image-caption pairs. For calibration,

we used the validation split, which comprises 8,738 unique image-caption pairs.

- **Nebula** (Matsuda et al., 2024): The Nebula dataset comprises 32,978 human annotations on image-caption pairs, spanning 32,978 unique images. It is organized into 26,382 training pairs, 3,298 validation pairs, and 3,298 test pairs. The dataset was introduced for training hallucination-robust image captioning evaluation metrics and was constructed by extending the Polaris dataset to increase visual diversity. In our experiments, we used the test split of Nebula, which contains 3,298 image-caption pairs. For calibration, we used the validation split, which also comprises 3,298 image-caption pairs.
- **EvalMuse-40K** (Han et al., 2024): The EvalMuse-40K dataset comprises 32,717 human annotations on image-caption pairs, spanning 32,717 unique images. It is organized into 29,445 training pairs and 3,272 validation pairs. EvalMuse-40K was introduced as a fine-grained benchmark for text-to-image generation evaluation, with comprehensive human annotations designed to better assess image-text alignment. Our experiments used the training pairs for model training, and the validation pairs for calibration.

Datasets used only for evaluation.

- **Flickr8K-Expert** (Hodosh et al., 2013): This dataset comprises 16,992 expert human judgments for 5,664 image–caption pairs from the Flickr8K dataset. Human assessors graded captions on a scale of 1 to 4, where 4 indicates a caption that accurately describes the image without errors, and 1 signifies a caption unrelated to the image.
- **Composite** (Aditya et al., 2015): This dataset contains 13,146 image–caption pairs taken from MS-COCO (2,007 images), Flickr8K (997 images), and Flickr30K (991 images). Each image originally had five reference captions. One of these references was chosen for human rating and subsequently removed from the reference set that is to be used when assessing evaluation metrics.

Table 3 summarizes the key statistics of the datasets used for training, calibration, and evaluation.

Additional Datasets.

The following datasets were used in our experiments to evaluate the proposed uncertainty estimator under different conditions, including controlled input ambiguity and model knowledge gaps. Specifically, we rely on them to assess calibration, sharpness, and correlation with human judgments, as well as to study how uncertainty changes under out-of-domain settings and varying caption lengths. IIW, DOCCI, Localized Narratives, and Docent are all datasets with long, detailed captions, making them particularly suitable for analyzing the behavior of the model in the long-caption setting.

IIW (Garg et al., 2024): ImageInWords (IIW) is a human-in-the-loop dataset designed to support rich, detailed image descriptions. It was created to capture hyper-detailed visual semantics and to facilitate the evaluation of models that generate long-form captions with fine-grained content. In our experiments, we use IIW as an in-domain dataset with realistic natural imagery and detailed textual descriptions, making it suitable for studying model behavior under standard captioning conditions. We evaluate results on the complete set of 400 test instances provided by the dataset.

DOCCI (Onoe et al., 2024): Descriptions of Connected and Contrasting Images (DOCCI) is a dataset of long, human-authored descriptions for approximately 15k images. The annotations are designed to be comprehensive and compositional, with an emphasis on challenges such as spatial relations, counting, text rendering, and world knowledge. We use DOCCI as an in-domain benchmark for long-caption evaluation, since it provides natural images paired with detailed descriptions that stress the model's ability to assess rich semantic alignment. In our experiments, we use the 100 test instances available in the dataset.

Localized Narratives (Pont-Tuset et al., 2020): Localized Narratives is a large-scale multimodal dataset in which annotators describe images while simultaneously hovering over the regions they are referring to, producing word-level grounding information. The dataset covers hundreds of thousands of images from sources such as COCO, Flickr30k, ADE20K, and Open Images, and is widely used for dense vision-language understanding. In our experiments, we use it as an in-domain dataset because it contains realistic scenes and natural language descriptions that are well aligned with standard image captioning evaluation settings. Since the test split

Metric	VICR	Polaris	Nebula	Flickr8K-Expert	Composite	EvalMuse-40K
Training	10175	26340	26382	-	-	29445
Validation	2310	8738	3298	-	-	3272
Testing	3161	8726	3298	5664	13146	-
Unique Images	9990	13691	32978	1000	3995	32717
Caption Len. $\mu \pm \sigma$ (Q1/Q2/Q3)	10.9 ± 3.0 (9.0/10.0/12.0)	9.3 ± 5.5 (8.0/10.0/11.0)	8.8 ± 2.5 (8.0/9.0/10.0)	11.9 ± 4.2 (9.0/11.0/14.0)	13.1 ± 9.3 (9.0/11.0/14.0)	11.0 ± 7.0 (6.0/9.0/14.0)
Annot./Inst. $\mu \pm \sigma$ (Q1/Q2/Q3)	4.4 ± 0.8 (4.0/4.0/5.0)	3.0 ± 2.8 (1.0/1.0/6.0)	1.0 ± 0.0 (1.0/1.0/1.0)	3.0 ± 0.0 (3.0/3.0/3.0)	1.0 ± 0.0 (1.0/1.0/1.0)	1.0 ± 0.0 (1.0/1.0/1.0)
Human Score $\mu \pm \sigma$ (Q1/Q2/Q3)	2.67 ± 1.41 (1.20/2.40/4.19)	0.25 ± 0.34 (0.00/0.00/0.62)	0.52 ± 0.36 (0.00/0.60/0.78)	1.58 ± 0.72 (1.00/1.33/2.00)	3.00 ± 1.45 (1.50/3.00/4.50)	0.46 ± 0.25 (0.25/0.42/0.67)
Rating Scale	[1.0, 5.0]	[0.0, 1.0]	[0.0, 1.0]	[1.0, 4.0]	[1.0, 5.0]	[0.0, 1.0]
Score Distribution						

Table 3: Statistical characterization of the datasets. For each metric (i.e., Caption Length, Annotations per Instance, and Human Score), we report the mean μ and standard deviation σ, along with quartiles (Q1: 25th percentile, Q2: median, Q3: 75th percentile). Score distributions are represented using sparkline histograms computed over the average human scores using the observed annotation scale inferred directly from each dataset.

contains more than 100k instances, we randomly sample 3k examples for our analysis.

MSCOCO (Lin et al., 2014): The Microsoft COCO Captions dataset is one of the most widely used benchmarks for image captioning, containing multiple human-generated captions for each image. It provides natural photographs with diverse everyday scenes and serves as a strong in-domain reference for both caption generation and evaluation. We use MSCOCO as an in-domain dataset in our experiments, including a random subset matched in size to ArtCaps for the short-caption setting. Specifically, we randomly sample 18k MSCOCO instances to compare directly against ArtCaps.

Docent (Ananthram et al., 2025): Docent is an artwork-oriented dataset used to evaluate captioning models under domain shift. Unlike natural-image benchmarks, it focuses on artistic imagery, where semantic interpretation can be more abstract and stylistically driven. We use Docent as an out-of-domain dataset to test whether the proposed uncertainty estimator reflects model knowledge gaps when processing artwork rather than photographic content. The dataset contains 100 distinct image-caption pairs, and we use the provided reference captions for our analysis.

ArtCaps (Lu et al., 2022): ArtCaps is a fine-art painting captioning dataset containing paintings paired with multiple human-written descriptions. It was constructed specifically for image captioning in the artistic domain, where visual content is often complex and culturally contextual. We use ArtCaps as an out-of-domain benchmark in our experiments, since it provides a clear contrast to natural-image datasets and is well suited for probing epistemic uncertainty under domain shift. The dataset contains approximately 18k instances, which we compare against a randomly sampled subset of 18k MSCOCO instances in the short-caption setting.

B Sensitivity to the Number of Calibration Clusters

We recall that the value $K = 63$ reported in the paper was not tuned on any coverage or width criterion. It was obtained from the elbow and silhouette curves of the K-means fit on the calibration image embeddings, that is, from an unsupervised criterion computed before any conformal procedure is run. This appendix reports a sensitivity analysis over the coverage and mean interval width over different test sets across different degrees of clustering.

For each value of K we cluster the validation split of the merged EvalMuse, Nebula, Polaris and VICR data, run the conformal prediction procedure at a target coverage of 90% inside every cluster to obtain a scaling factor λ and an additive bias b, and retain the most conservative pair across clusters. Those parameters are applied to the original uncalibrated quantiles of each benchmark, and we report the resulting empirical coverage and mean interval width. The point $K = 1$ is the unclustered calibration and serves as a baseline.

Figure 5 shows that coverage and width move together with K. The parameters applied at test

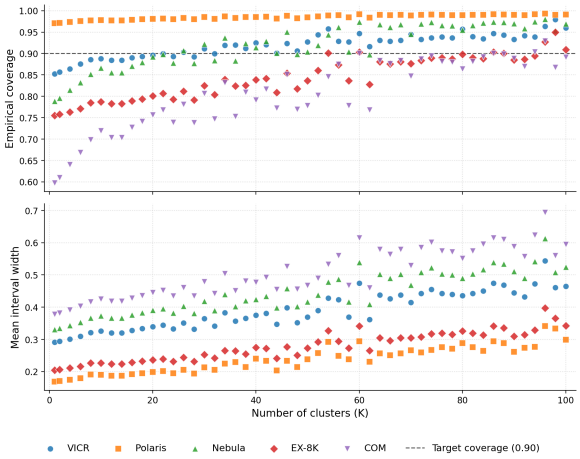


Figure 5: Empirical coverage (top) and mean interval width (bottom) as a function of the number of calibration clusters K. Colours identify the benchmark and are shared between the two panels. The dashed line marks the target coverage of 0.90 and the dash-dotted line marks the value $K = 63$ used in the paper. The point at $K = 1$ corresponds to calibrating on the whole validation split without clustering.

time are the most conservative pair found across clusters, so increasing K increases the number of candidates over which that maximum is taken, and the widening is attributable to the selection rule rather than to a finer conditioning on image content.

C Input Ambiguity Beyond Blur

While Gaussian blur provides a controlled way to degrade visual information, it represents only one form of input ambiguity. In practice, images may be affected by a wide range of distortions that alter different aspects of the visual signal, such as noise, compression artifacts, structural occlusions, or geometric deformations. To assess whether our conclusions for BEiTScore++ generalize beyond blur, we extend the analysis to additional perturbations, including salt-and-pepper noise, JPEG compression, random patch cutout, contrast adjustments, and elastic warping. Figure 7 presents a grid of examples illustrating these different distortions.

We apply each distortion to the images and measure the corresponding uncertainty estimates. The resulting distributions of confidence interval widths are presented in Figure 6. For most distortion types, we observe that introducing degradation leads to increased uncertainty compared to the no-distortion setting. This trend is consistent with the results obtained using Gaussian blur, and further supports the hypothesis that the proposed estimator is sensitive to input ambiguity.

An exception is observed for contrast adjustments, where the interval widths remain unchanged. We find similar behavior when applying other color-based transformations, suggesting that the model is relatively robust to photometric changes that preserve the underlying geometry and shape of the image content. In contrast, distortions that disrupt structural or spatial information (e.g., noise, occlusions, and warping) lead to a clear increase in model uncertainty.

D Analysis of Different Sources of Uncertainty Beyond BEITScore-L

In this appendix, we report the results obtained with the base-sized version of BEiTScore++ for the same set of experiments described in Section 5.4. Overall, the base model exhibits the same behavior as the large model, confirming that the observed trends are not specific to model scale. In particular, the uncertainty estimates increase consistently with stronger input corruptions, remain stable under color-based perturbations, and rise on out-of-domain artwork datasets, in comparison to the in-domain natural image benchmarks, due to model knowledge gaps.

D.1 Input Ambiguity

We first analyze the effect of input ambiguity under Gaussian blur. Some visual examples of the instance images, with different blur levels applied to them, can be seen in Figure 8. We use the Polaris test split for this analysis and progressively increase the blur strength applied to each image. To reduce the confounding effect of the highly skewed score distribution in Polaris, we again downsample instances with a human score of zero, obtaining a more balanced set of image-caption pairs. This allows us to more clearly examine how uncertainty changes as the visual information becomes progressively degraded.

Figure 9 shows the confidence interval widths produced by the base model under increasing blur levels. The results follow the same pattern observed for the large model: the interval width grows with blur intensity, indicating that the estimator responds appropriately to increasing input ambiguity. As before, the strongest blur levels produce a slight reduction in interval width, which we attribute to the model becoming confident that the image-caption pair is strongly misaligned.

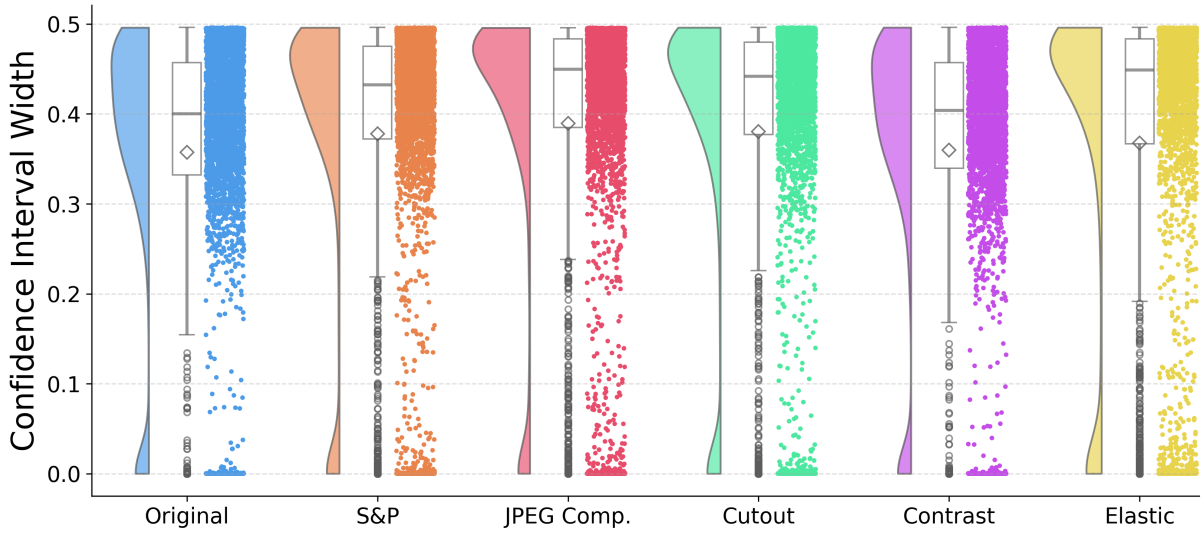


Figure 6: Box plots of confidence interval widths produced by our uncertainty estimator, associated to the BEiTScore++ L/14 model, under different types of image distortions.

We also evaluate the base model under the broader set of distortions described in the previous Section C, including salt-and-pepper noise, JPEG compression, random patch cutout, elastic warping, and contrast adjustments. The results are consistent with the large model, and can be seen in Figure 10. Most distortions lead to larger confidence intervals as their severity increases, while contrast and other color-based transformations produce little to no change in uncertainty. This suggests that the base model is likewise robust to purely photometric perturbations and primarily sensitive to structural corruption of the image.

D.2 Model Knowledge Gaps

We next examine whether the base model reflects uncertainty arising from model knowledge gaps. We compare in-domain datasets containing natural images and realistic scenes, namely IIW, DOCCI, Localized Narratives, and MSCOCO, against out-of-domain artwork datasets, namely Docent and ArtCaps. We again split the evaluation into long-caption and short-caption settings, to match the caption-length behavior expected by BEiTScore.

For long captions, we compare IIW, DOCCI, and Localized Narratives against Docent. For short captions, we compare ArtCaps against a randomly sampled subset of MSCOCO of equal size.

Figure 11 shows that the base model, like the large model, assigns substantially higher uncertainty to artwork datasets in both settings. This indicates that the uncertainty estimator captures epistemic uncertainty induced by domain shift, even at smaller model scale.

D.3 Marginal Distributions for the BEiTScore++ Predicted Quantiles

This section reports, for every evaluated dataset, the marginal distribution of each predicted quantile produced by BEiTScore++-L, together with the distribution of the mean human ratings. Each panel aggregates the prediction, prior to conformal calibration, over all instances in the corresponding dataset, thus, the histograms describe how a given quantile level is distributed across the evaluation set rather, than the interval predicted for any single caption. The distributions for the VICR, Polaris, Nebula, Ex-8K, and COM datasets are respectively shown in Figures 12 to 16.

The same qualitative behaviour is observed across all datasets. The panels display the expected monotone ordering of the quantile levels. The lower quantiles, from the 5th to the 12.5th, are heavily concentrated near the lower end of the score range, with the bulk of their mass in that region. As the quantile level increases, the mass progressively shifts to the right and spreads over the full support of the score. The 50th quantile, which corresponds to the median prediction, covers the entire range and its shape closely tracks the distribution of the mean human ratings shown in the same figure. The upper quantiles, from the 85th to the 95th, are shifted towards the higher end of the range and concentrate near the upper bound, mirroring the behaviour of the lower quantiles at the opposite extreme.

This pattern is consistent with the way the model constructs its prediction intervals. Because the quantile head produces nested intervals of decreasing nominal coverage, the lower and upper quantiles act as the floor and ceiling of these intervals and therefore accumulate near the extremes of the score range, while the median remains aligned with the target distribution. The absence of quantile crossing across the panels, together with the alignment between the median and the human ratings, indicates that the predicted quantiles are well ordered and that the central prediction is calibrated to the scale of the human judgements. These observations hold consistently across the five datasets, which indicates that the behaviour is a property of the model rather than an artefact of any particular evaluation set.

D.4 Qualitative Examples

This section presents qualitative examples of the predictions produced by BEiTScore++-L for every evaluated dataset. For each example the correesponding figure shows the input image, the candidate caption, the median score predicted by the model, and the 90 percent prediction interval obtained after applying the calibrated scaling and bias correction. The mean human rating is reported alongside each prediction so that the point estimate and the predicted interval can be compared directly against the ground truth. The examples were sampled at random from each evaluation set, without any selection based on the quality of the prediction, so that the figures reflect the typical behaviour of the model rather than a curated subset.

The examples illustrate the general behaviour discussed in the main text. In the majority of cases the mean human rating falls within the predicted 90 percent interval, which is consistent with the calibrated coverage reported for the model. The width of the interval also varies with the difficulty of the instance, with wider intervals appearing for captions whose agreement with the image is more ambiguous and narrower intervals for clear cases at the extremes of the score range. Cases in which the human rating falls outside the interval are also present, as expected for a method that targets a nominal coverage level rather than exact containment. These cases are distributed across the datasets without any obvious systematic pattern.

Multiple examples for the each of the VICR, Polaris, Nebula, Ex-8K, and COM datasets, are shown in Figures 17 to 26.

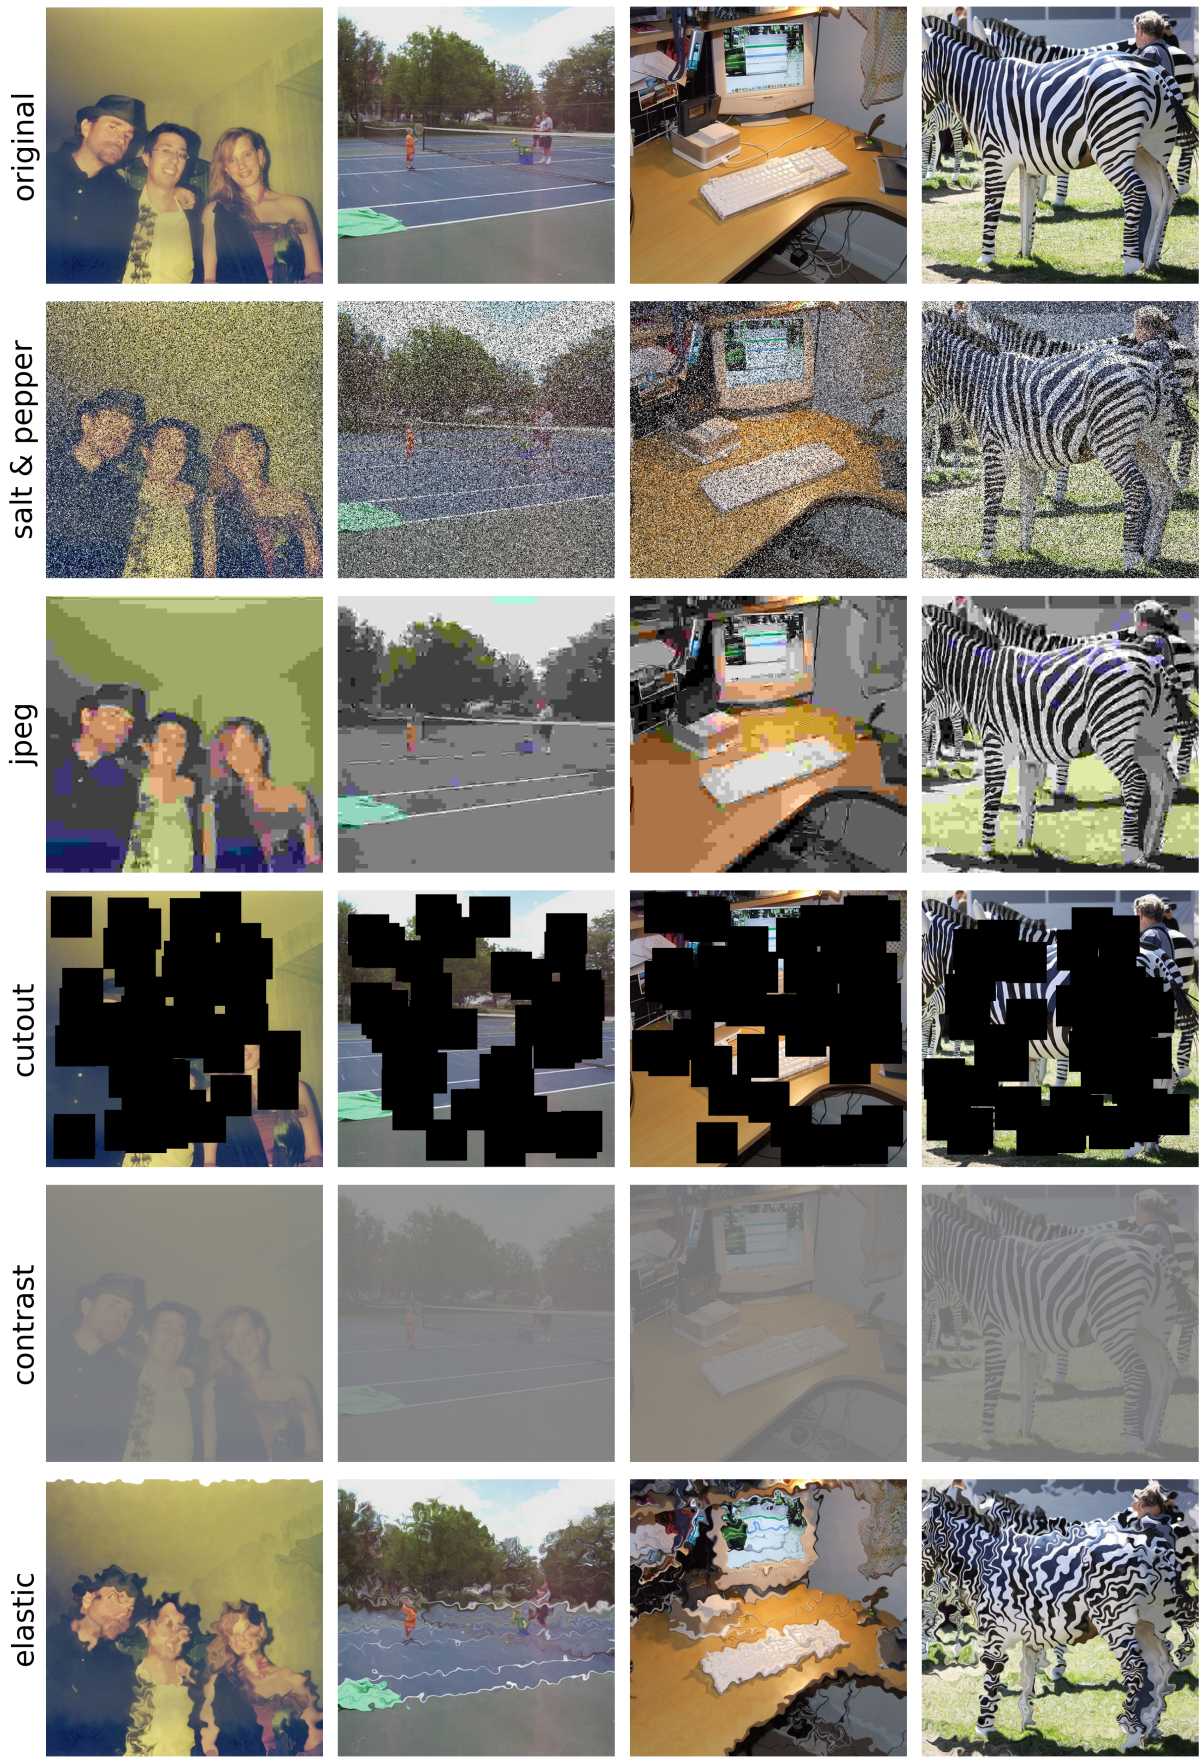


Figure 7: Examples of image distortions used to evaluate input ambiguity beyond Gaussian blur. Each row corresponds to a different type of perturbation.

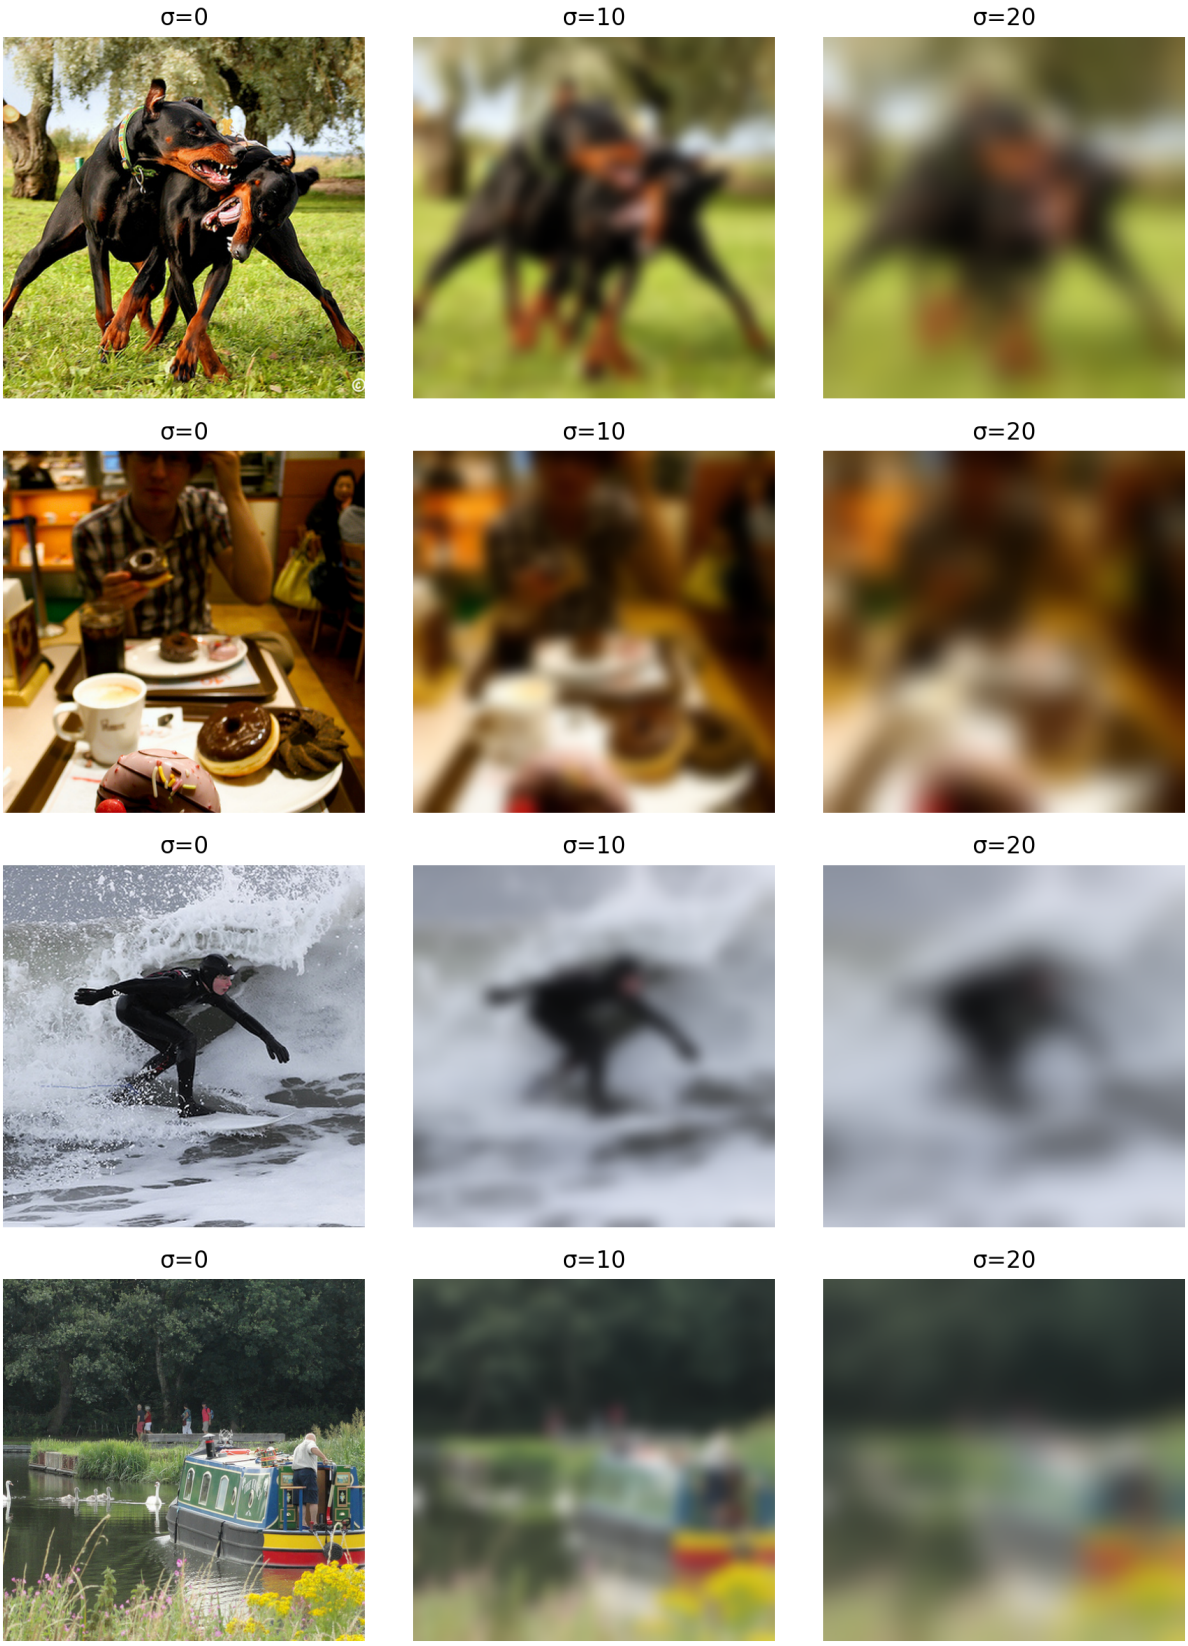


Figure 8: Examples showing different levels of blur applied to four images from the Polaris test split, used during the evaluation of input ambiguity.

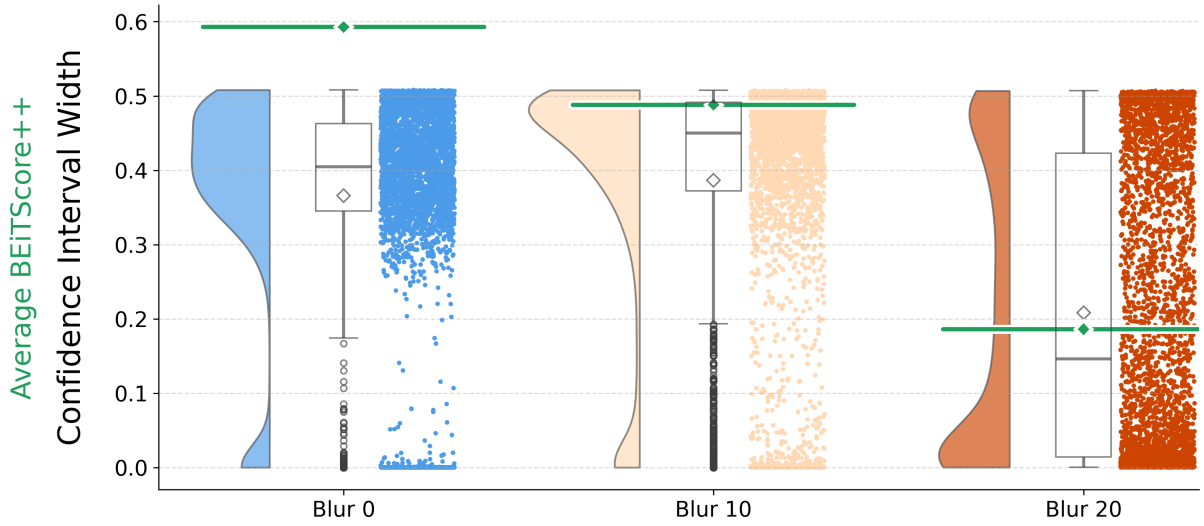


Figure 9: Distribution of 90% confidence interval widths generated by the base variant of BEiTScore++ as Gaussian blur increases on images from the Polaris test split. The horizontal lines indicate the average median score produced by BEiTScore++ at each blur level.

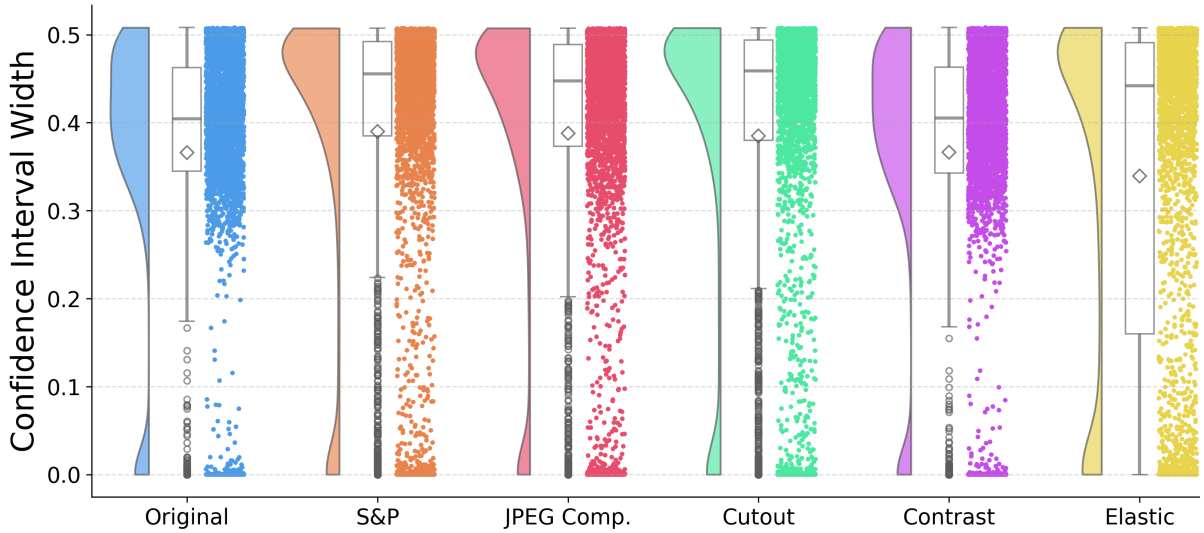


Figure 10: Box plots of confidence interval widths produced by the base-sized BEiTScore++ model, under different types of distortions applied to images from the Polaris test split.

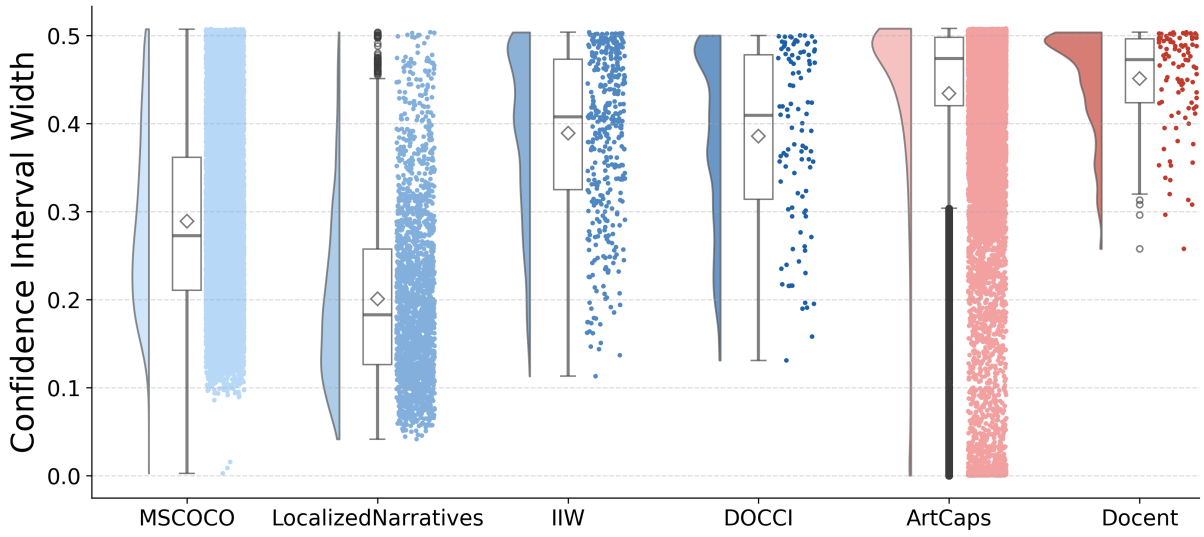


Figure 11: Distribution of 90% confidence interval widths produced by the base variant of BEiTScore++ for in-domain (natural images) and out-of-domain (artwork) datasets, across both short- and long-caption settings. The color gradient represents caption length, with pastel tones indicating short and direct captions, and with more saturated tones corresponding to datasets with more detailed descriptions.

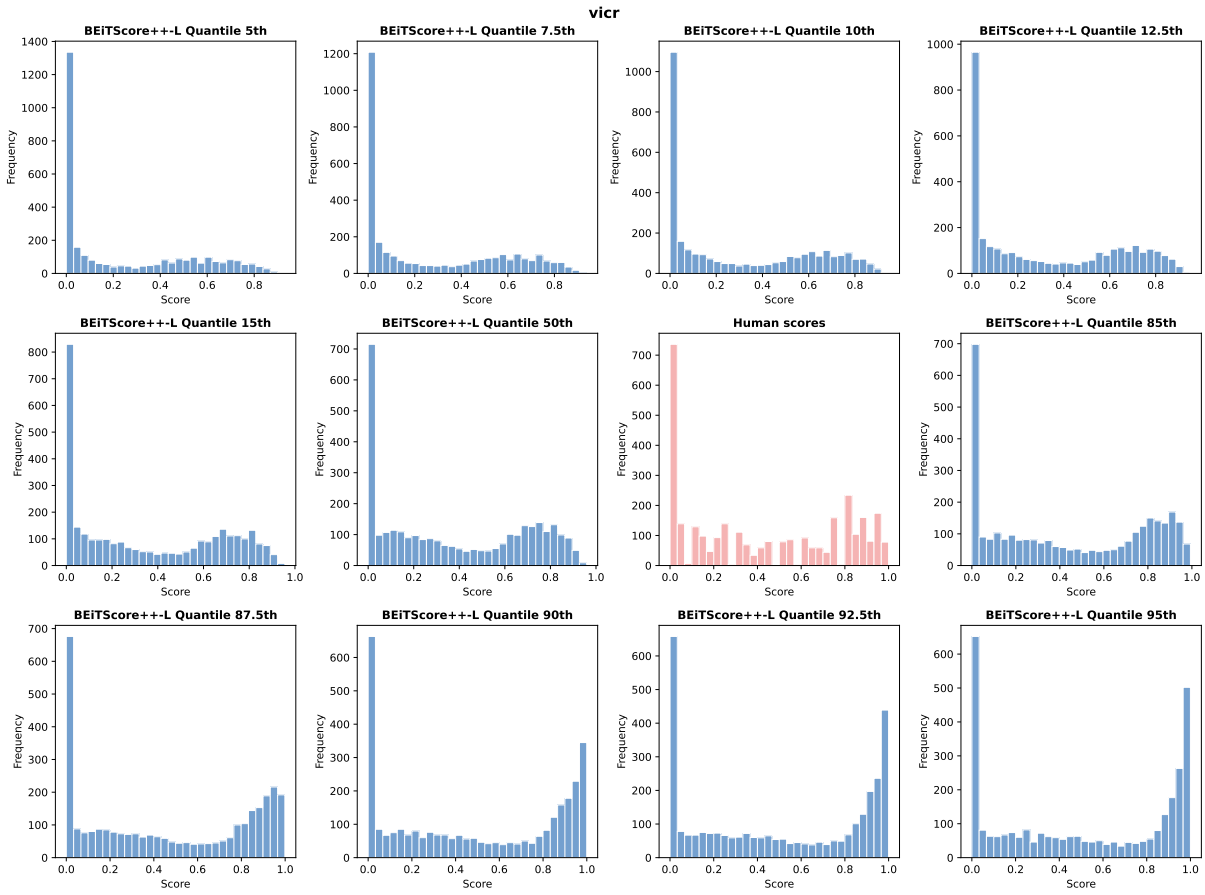


Figure 12: Marginal distributions of the predicted quantiles and the mean human ratings for the VICR dataset.

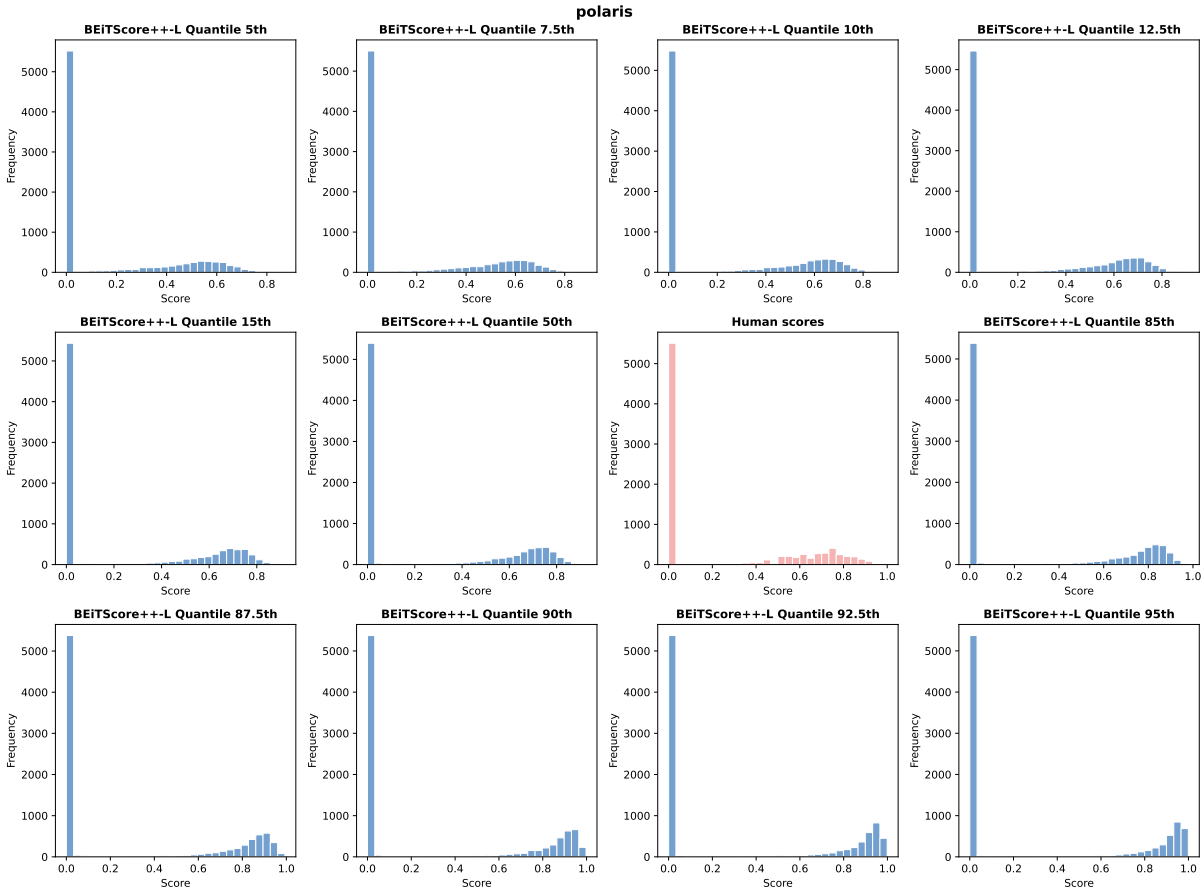


Figure 13: Marginal distributions of the predicted quantiles and the mean human ratings for the Polaris dataset.

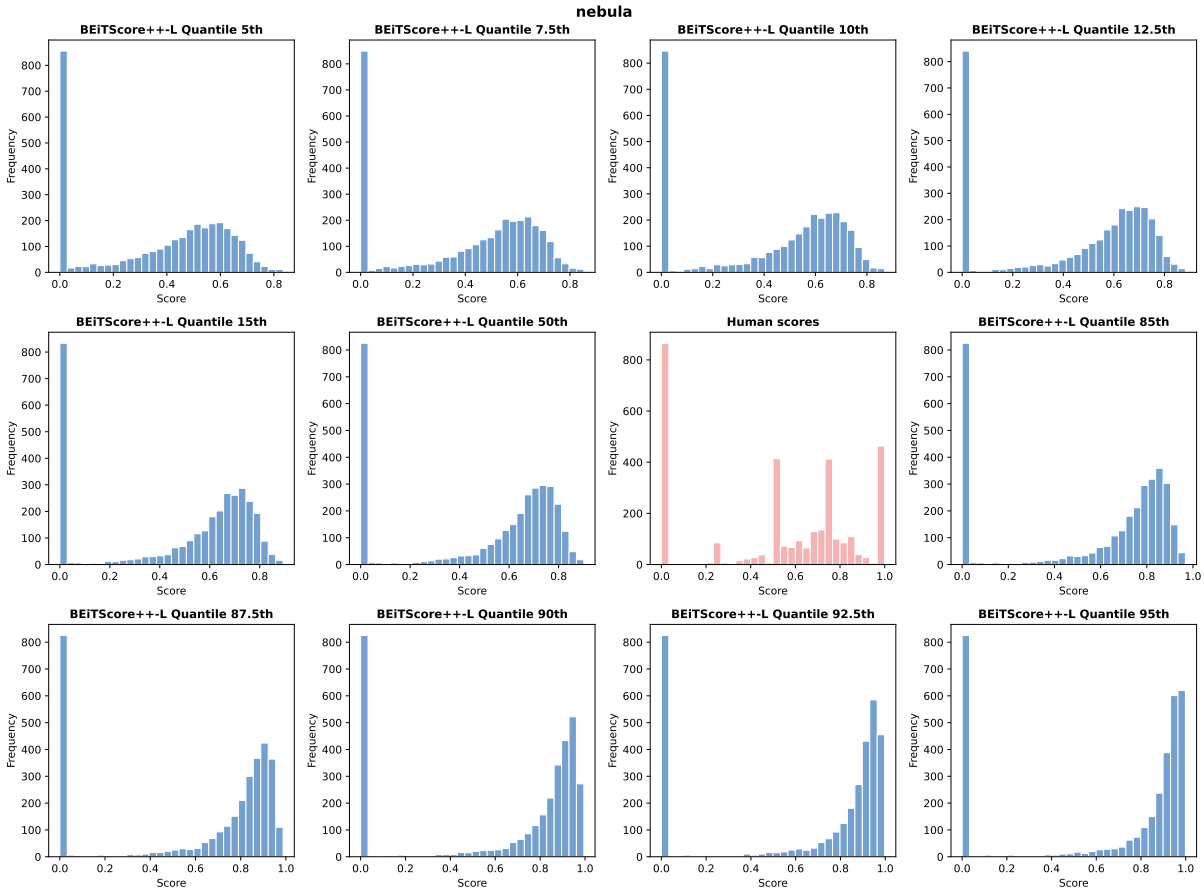


Figure 14: Marginal distributions of the predicted quantiles and the mean human ratings for the Nebula dataset.

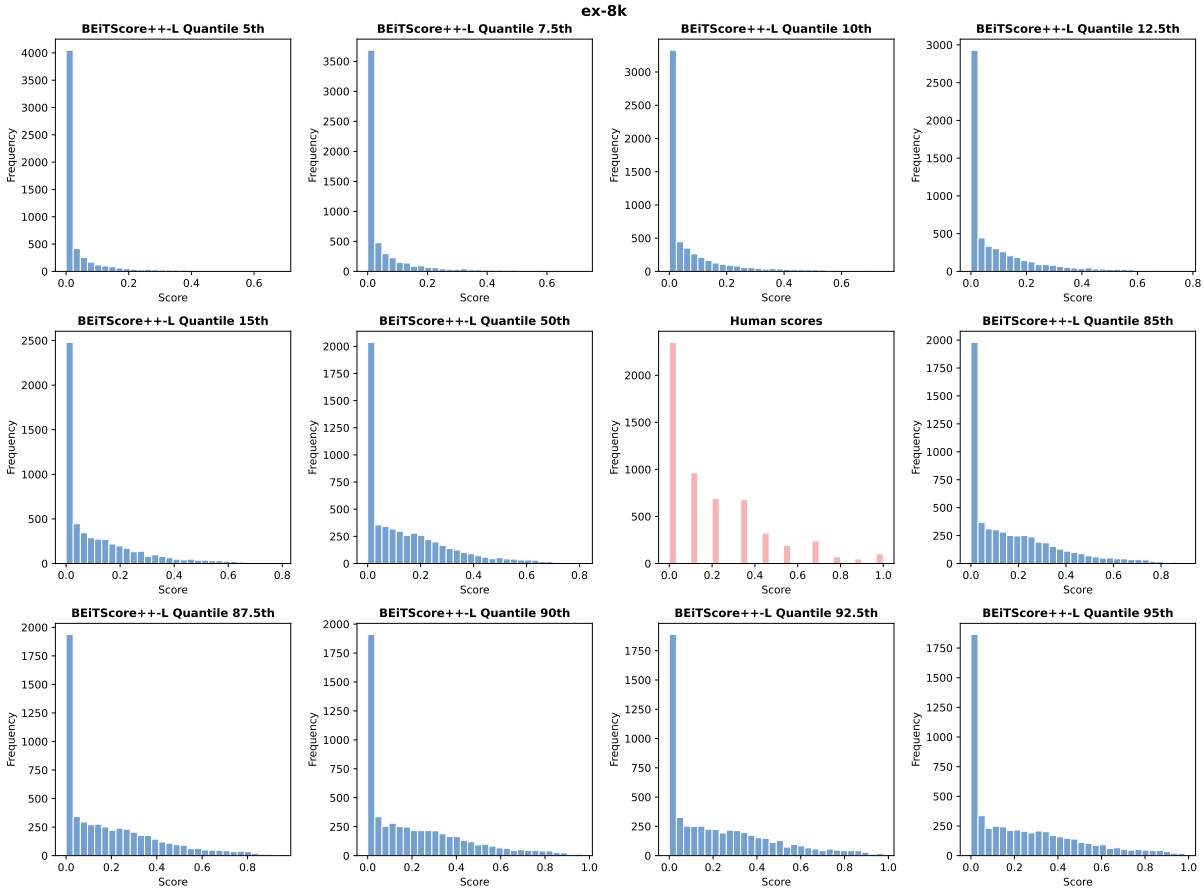


Figure 15: Marginal distributions of the predicted quantiles and the mean human ratings for the Ex-8K dataset.

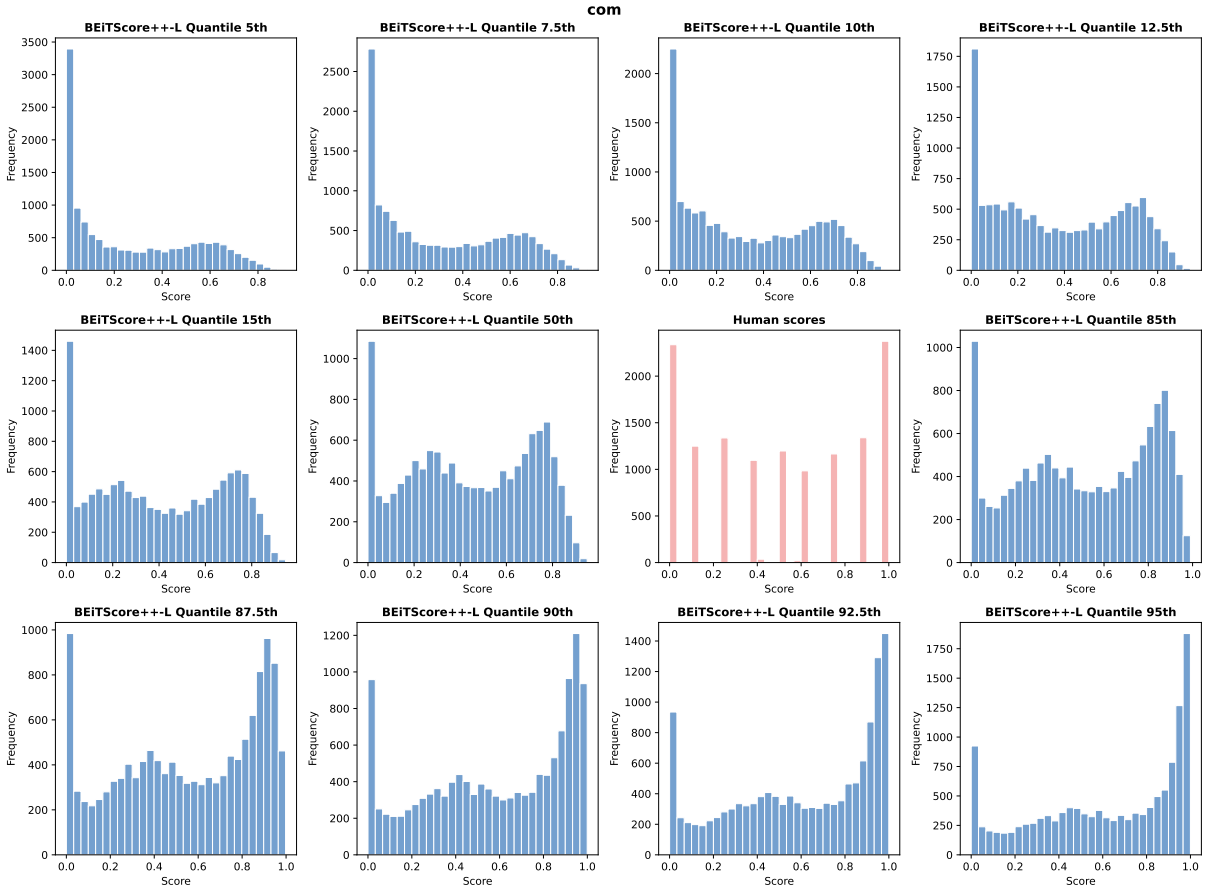


Figure 16: Marginal distributions of the predicted quantiles and the mean human ratings for the COM dataset.

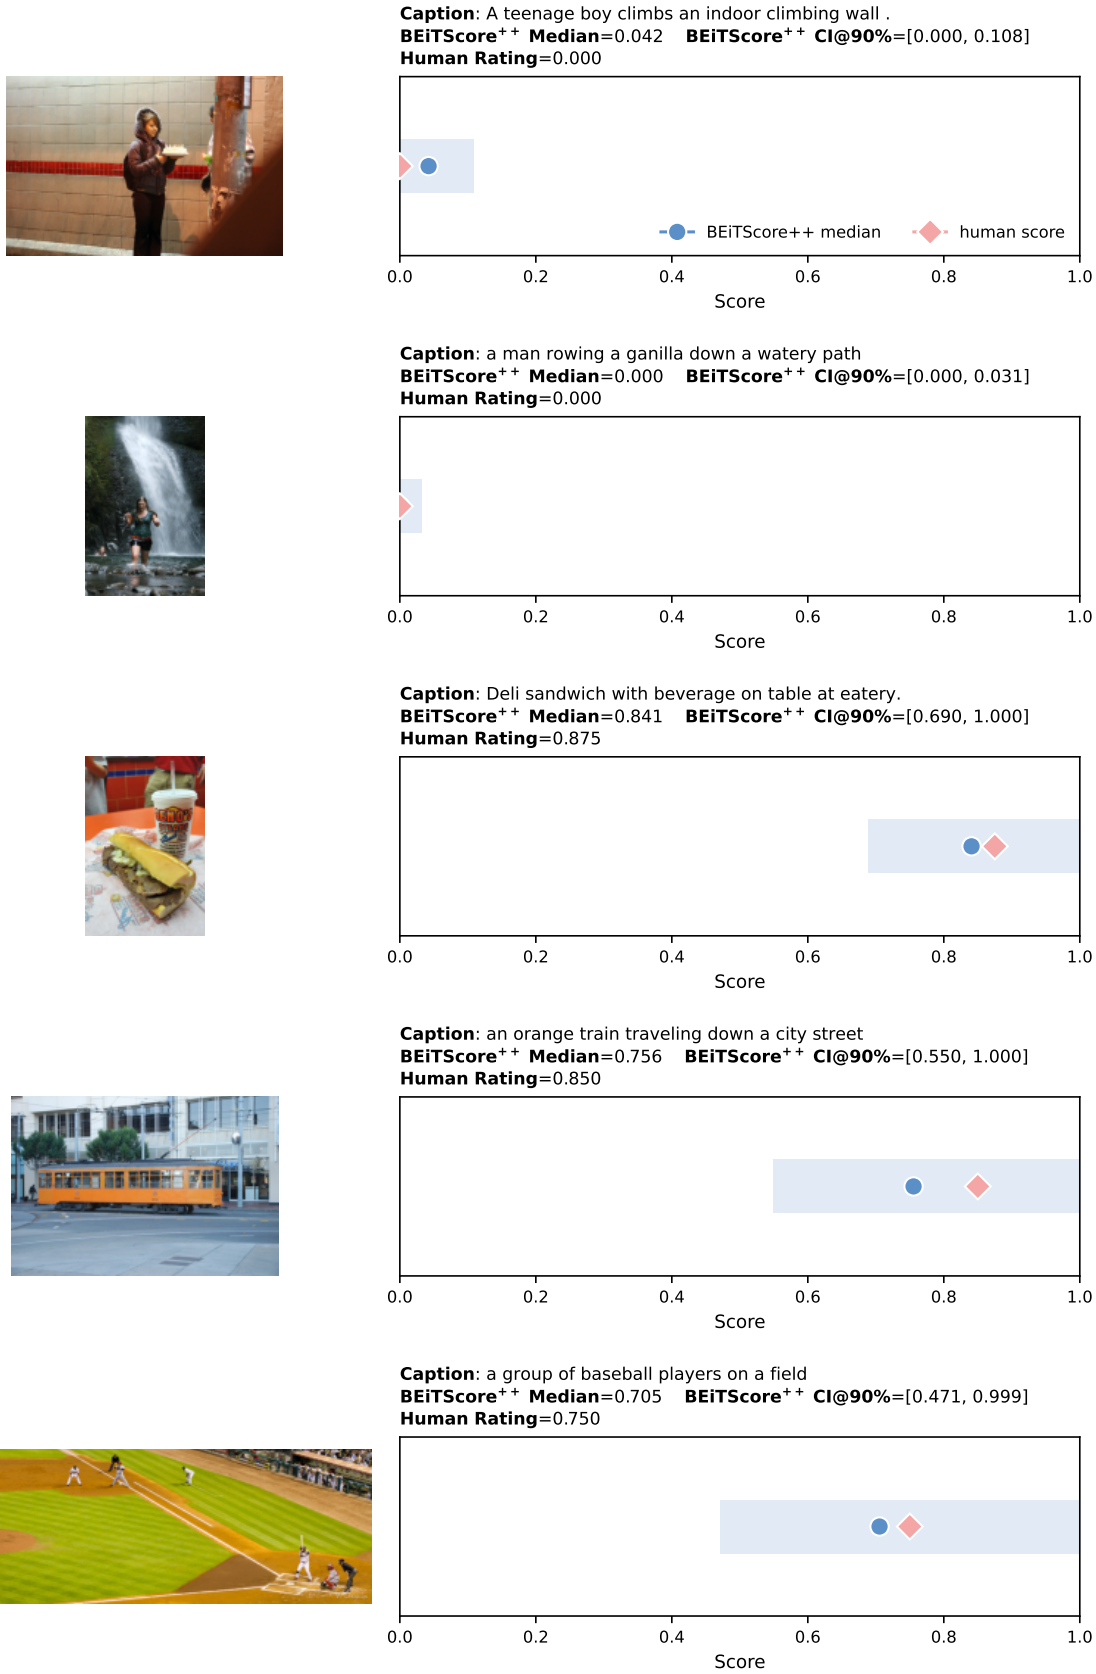
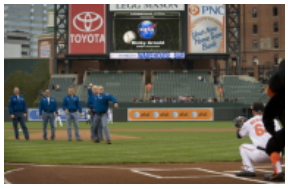
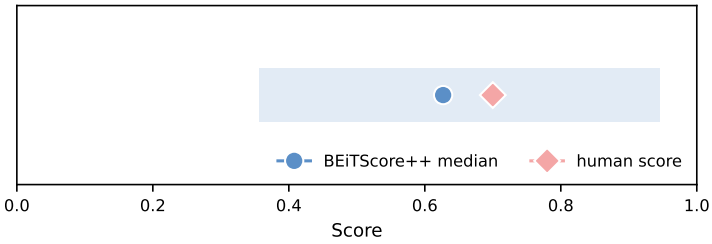


Figure 17: Qualitative examples for the predictions of BEiTScore++-L on the VICR dataset. Each row shows the input image, the candidate caption, the predicted median and 90 percent interval, and the mean human rating.

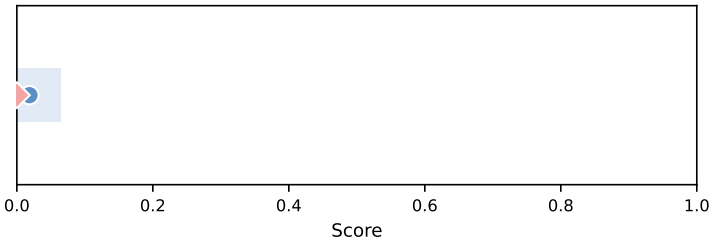


Caption: a group of men playing a game of baseball
BEiTScore^{++} Median=0.627 **BEiTScore^{++} CI@90%**=[0.356, 0.945]
Human Rating=0.700



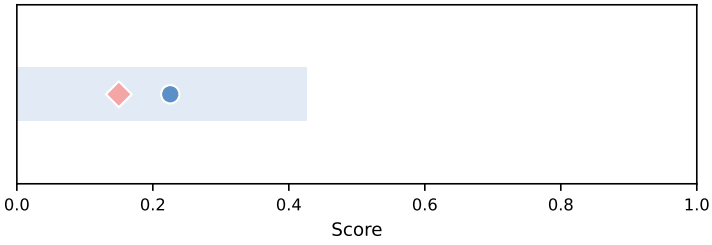


Caption: Two people sit on a large rock near a city street .
BEiTScore^{++} Median=0.019 **BEiTScore^{++} CI@90%**=[0.000, 0.065]
Human Rating=0.000



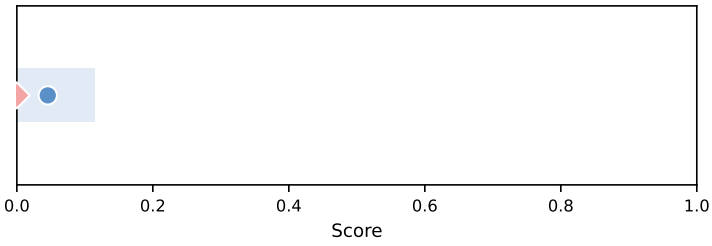


Caption: A woman next to a dog which is running an obstacle course .
BEiTScore^{++} Median=0.226 **BEiTScore^{++} CI@90%**=[0.000, 0.426]
Human Rating=0.150



Caption: A woman in a red coat is smiling , while people in the background are walking around in winter clothing .
BEiTScore^{++} Median=0.045 **BEiTScore^{++} CI@90%**=[0.000, 0.114]
Human Rating=0.000





Caption: A man in a bright green shirt riding a bicycle in a paved courtyard , carying a drink .
BEiTScore^{++} Median=0.166 **BEiTScore^{++} CI@90%**=[0.000, 0.327]
Human Rating=0.250



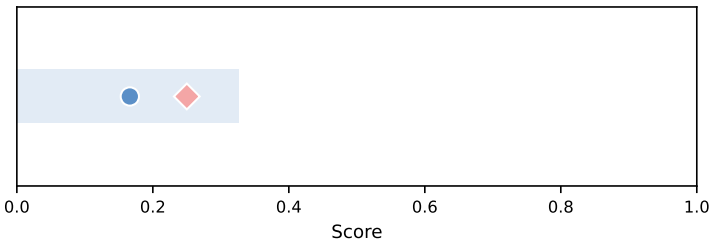


Figure 18: Additional qualitative examples for the predictions of BEiTScore++-L on the VICR dataset.

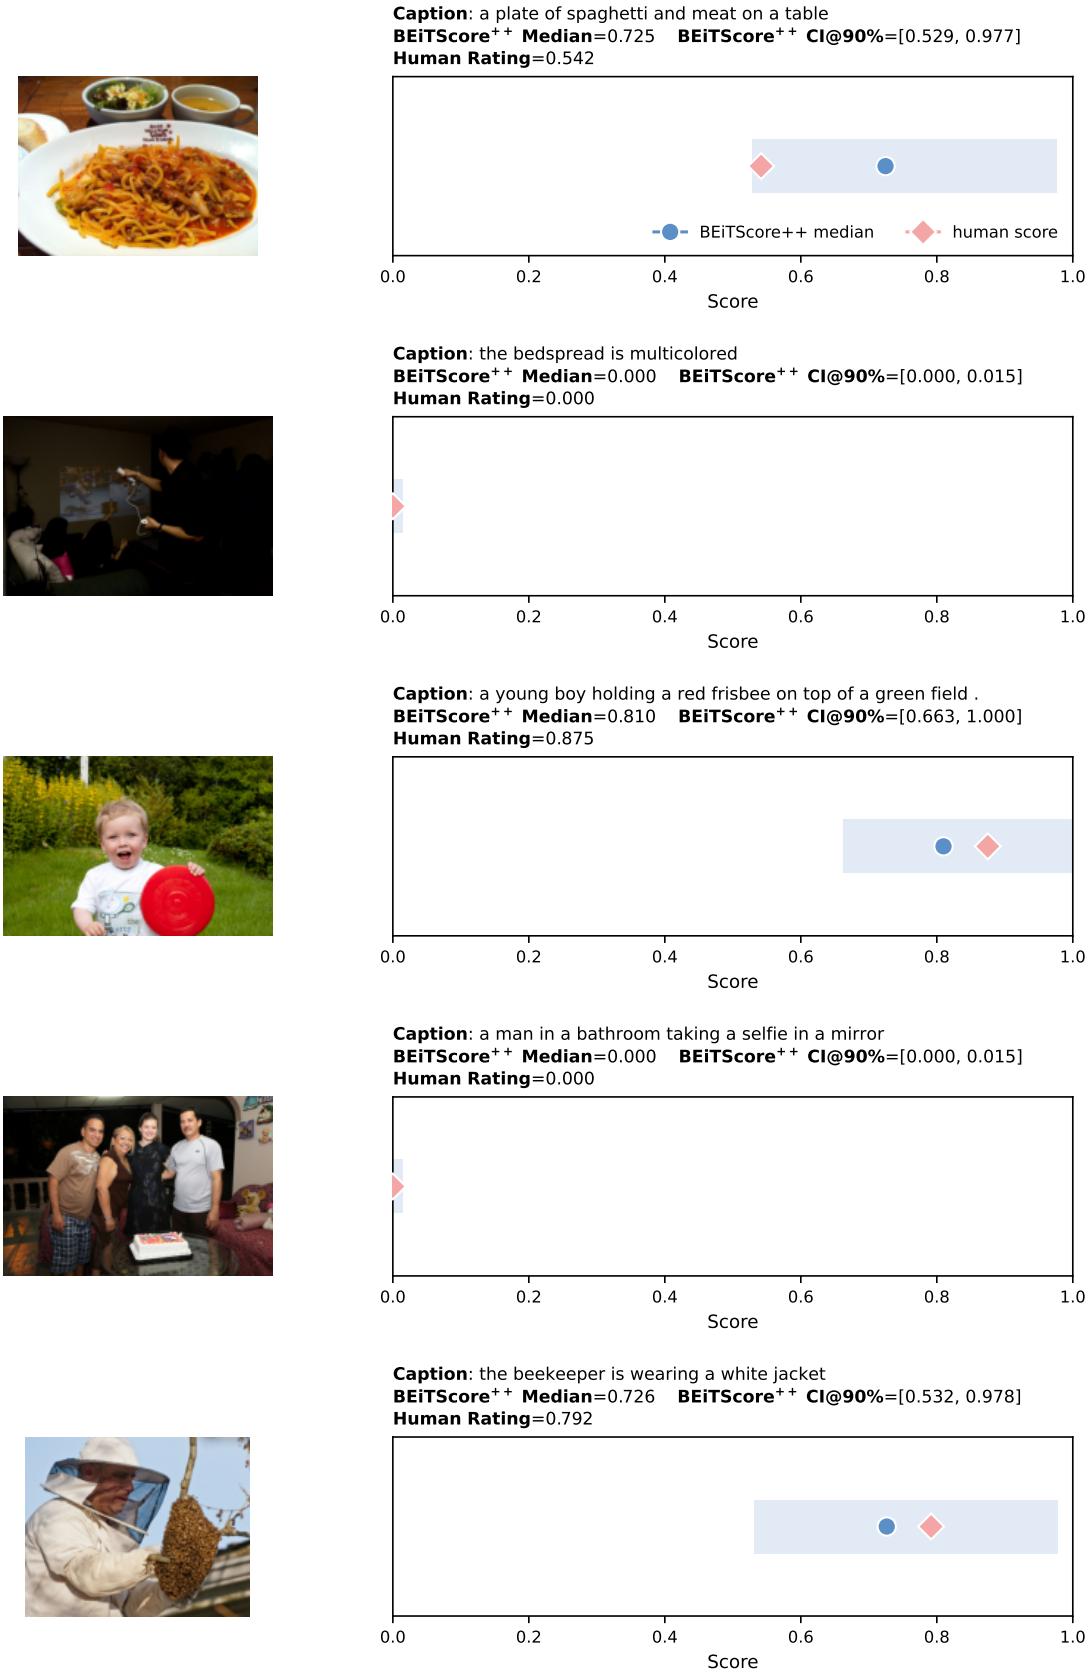
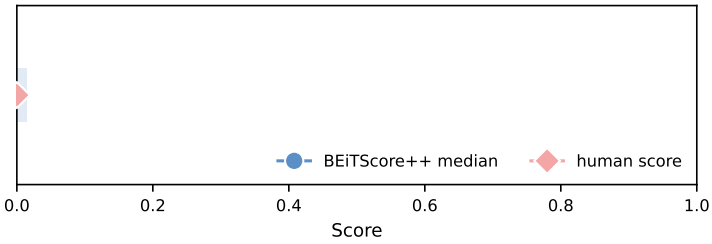


Figure 19: Qualitative examples for the predictions of BEiTScore++-L on the Polaris dataset. Each row shows the input image, the candidate caption, the predicted median and 90 percent interval, and the mean human rating.

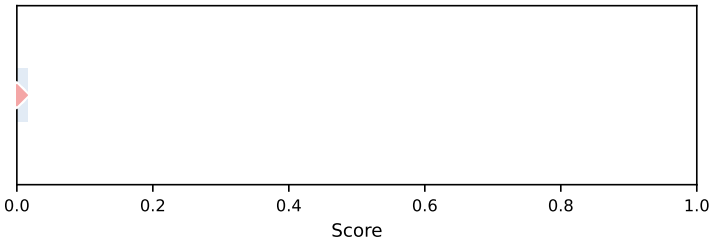


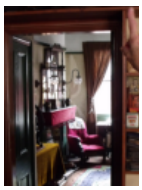
Caption: a pair of scissors and a blue bag on a table
BEiTScore^{++} Median=0.000 **BEiTScore^{++} CI@90%**=[0.000, 0.015]
Human Rating=0.000



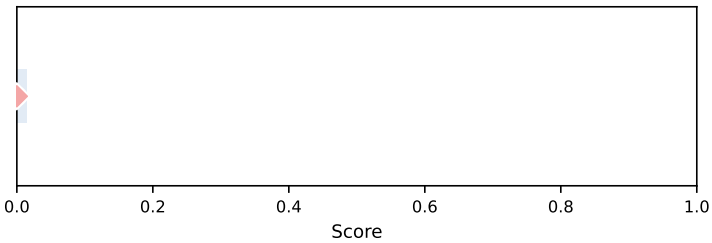


Caption: a couple of people standing on top of a snow covered mountain
BEiTScore^{++} Median=0.001 **BEiTScore^{++} CI@90%**=[0.000, 0.016]
Human Rating=0.000



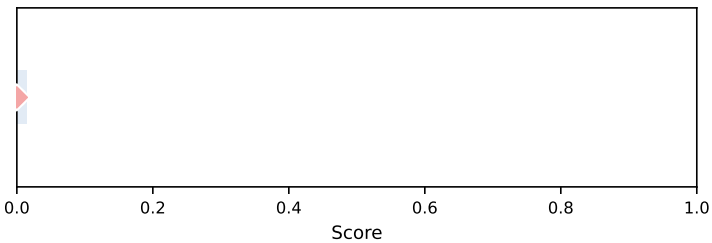


Caption: two pizzas sitting in pans on top of a table .
BEiTScore^{++} Median=0.000 **BEiTScore^{++} CI@90%**=[0.000, 0.015]
Human Rating=0.000





Caption: a group of people posing for a picture with a cake
BEiTScore^{++} Median=0.000 **BEiTScore^{++} CI@90%**=[0.000, 0.015]
Human Rating=0.000





Caption: a group of soldiers
BEiTScore^{++} Median=0.542 **BEiTScore^{++} CI@90%**=[0.276, 0.830]
Human Rating=0.536

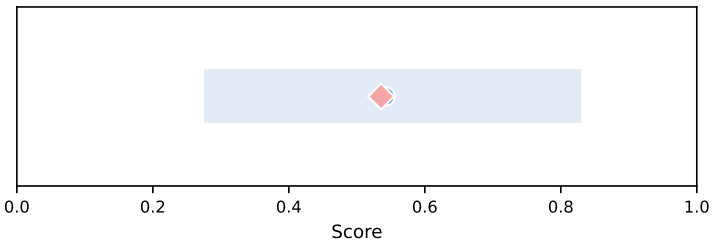


Figure 20: Additional qualitative examples for the predictions of BEiTScore++-L on the Polaris dataset.

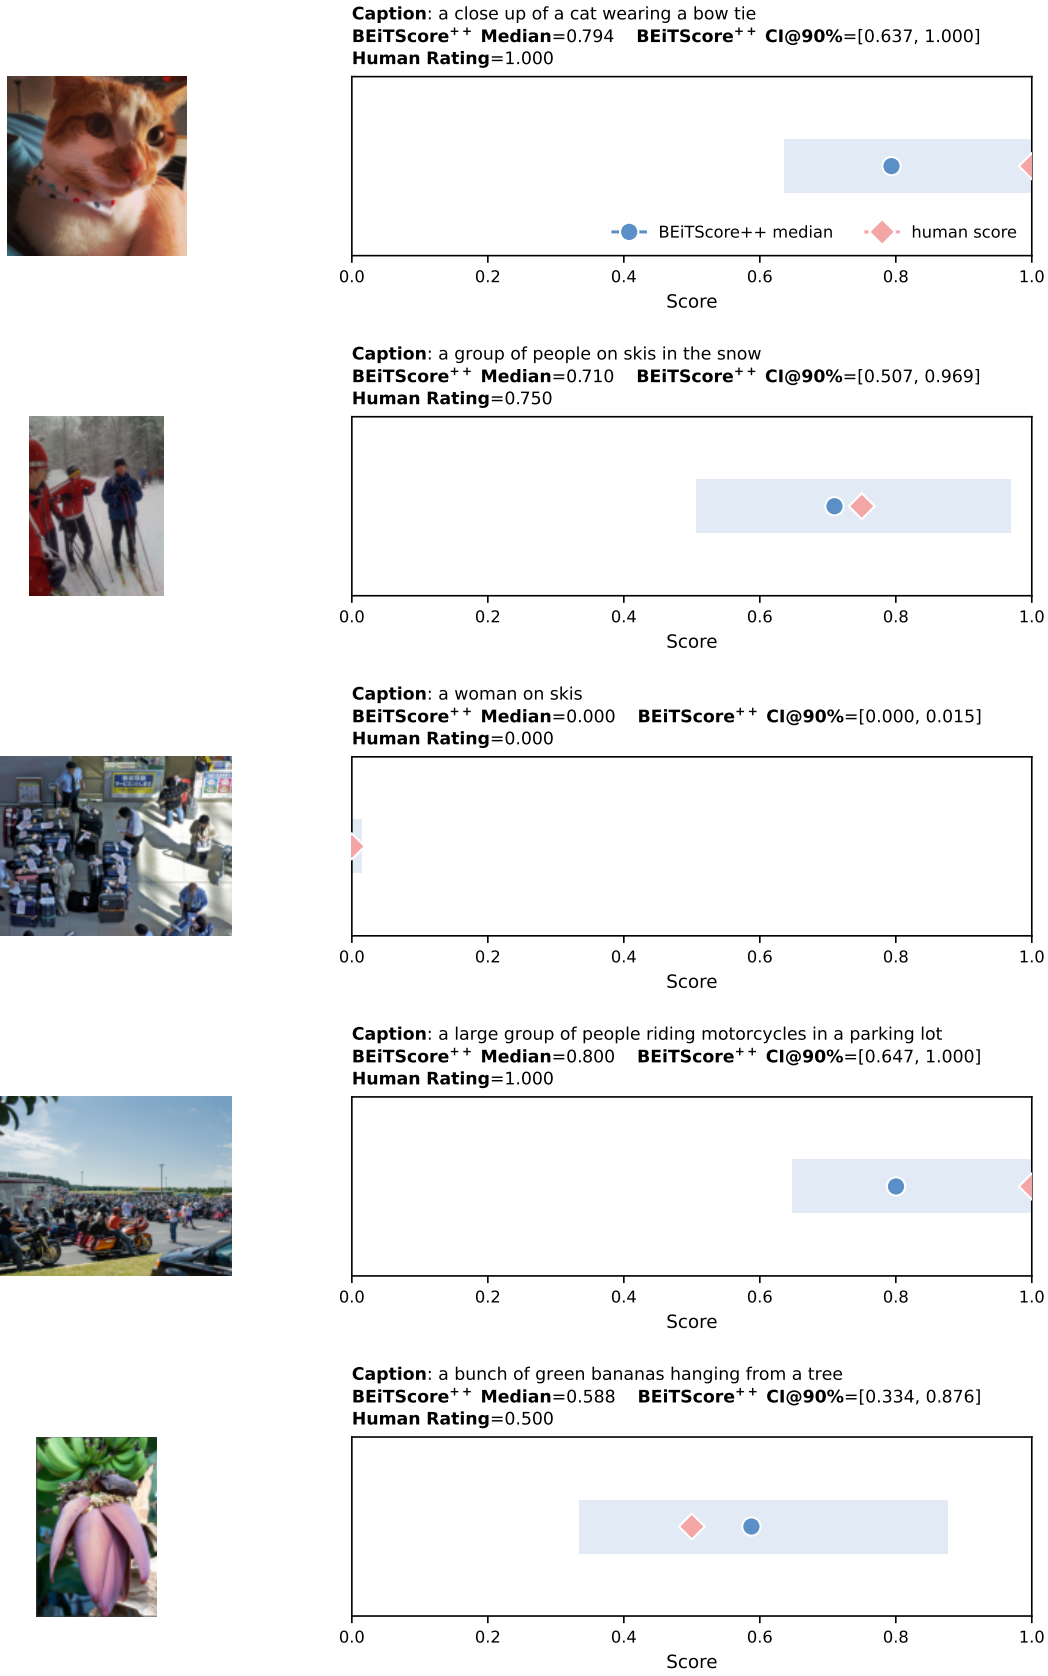


Figure 21: Qualitative examples for the predictions of BEiTScore++-L on the Nebula dataset. Each row shows the input image, the candidate caption, the predicted median and 90 percent interval, and the mean human rating.

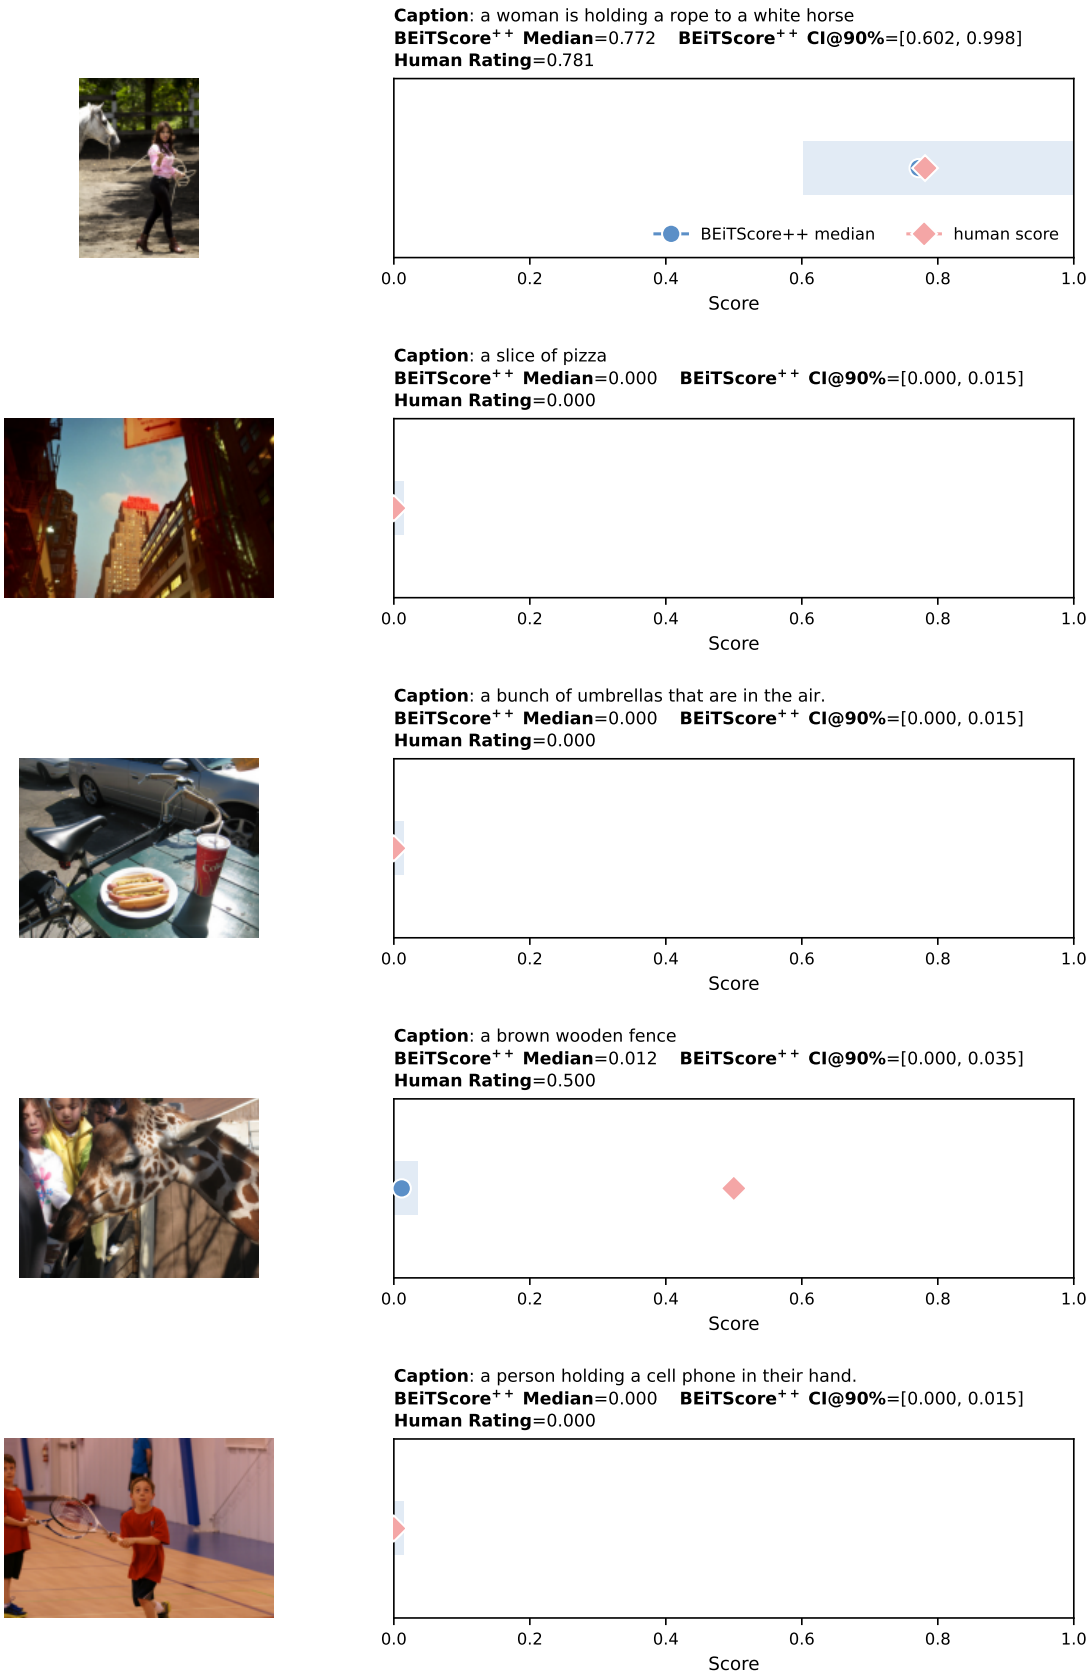


Figure 22: Additional qualitative examples for the predictions of BEiTScore++-L on the Nebula dataset.

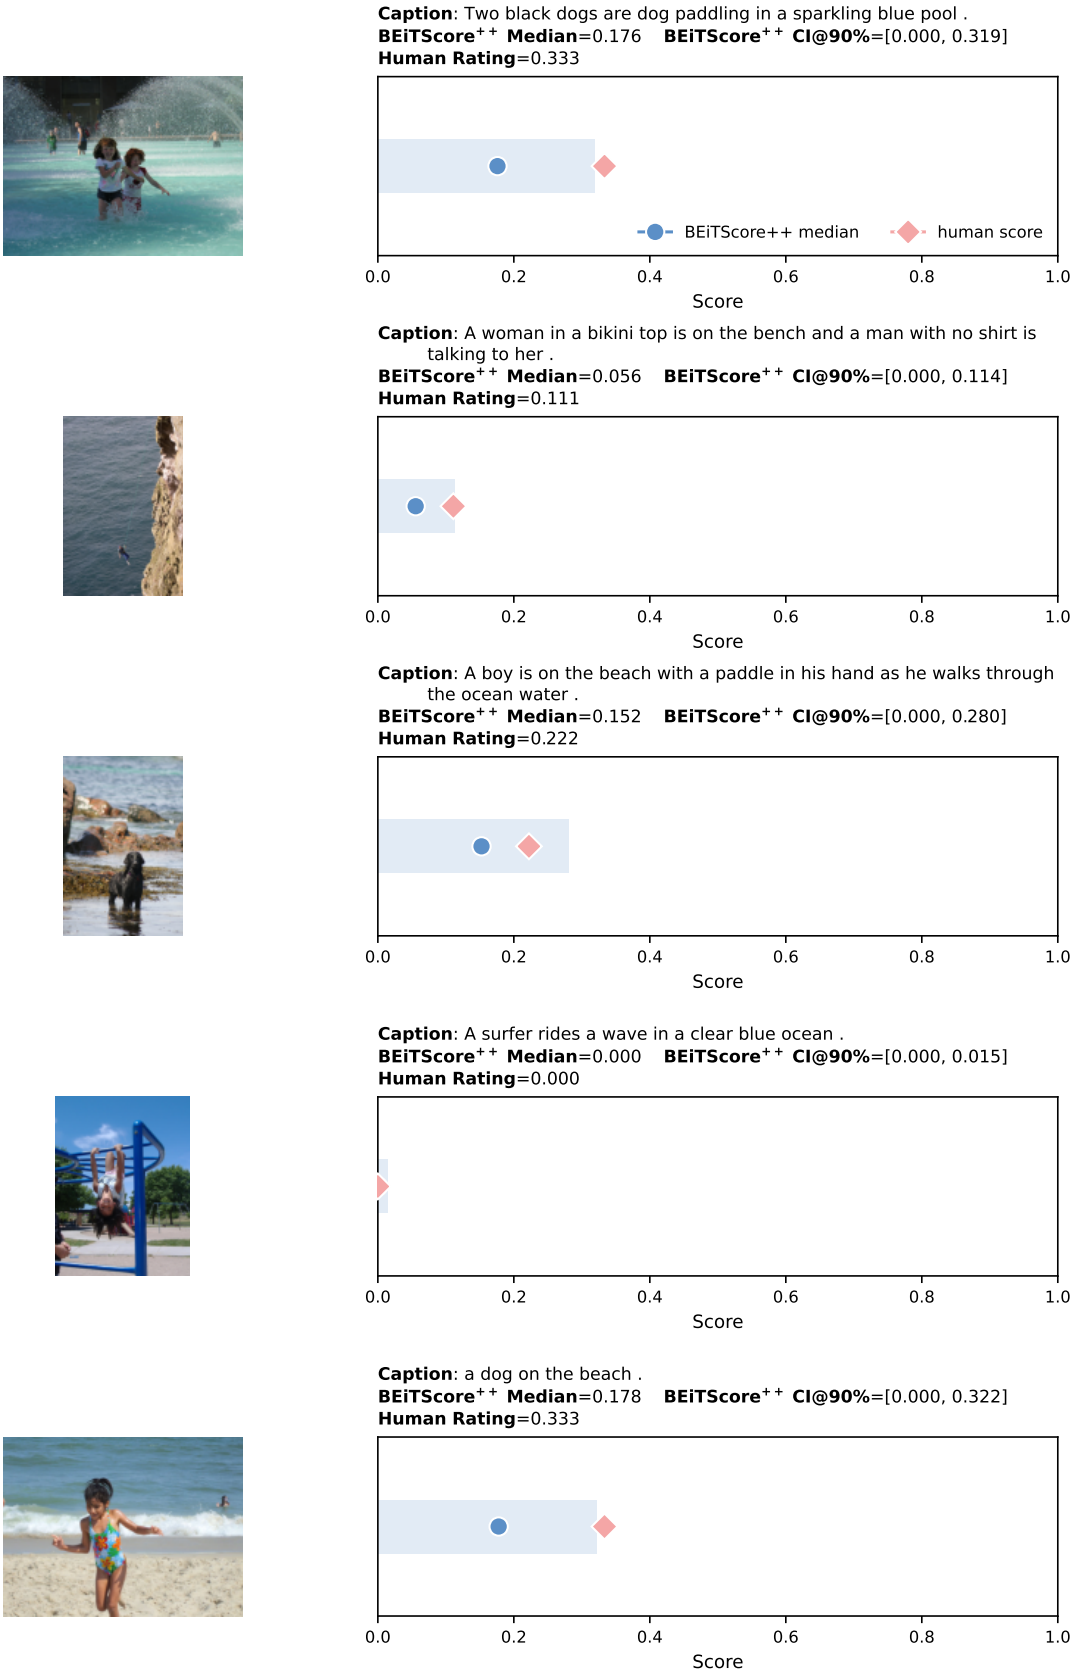
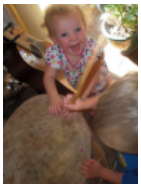
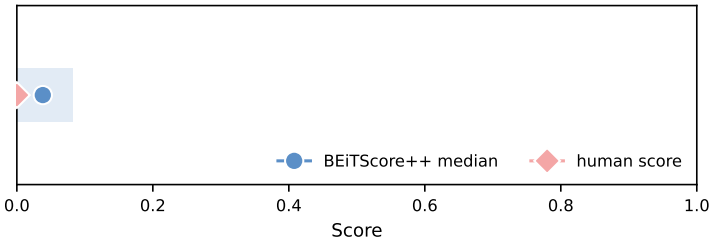


Figure 23: Qualitative examples for the predictions of BEiTScore++-L on the Ex-8K dataset. Each row shows the input image, the candidate caption, the predicted median and 90 percent interval, and the mean human rating.

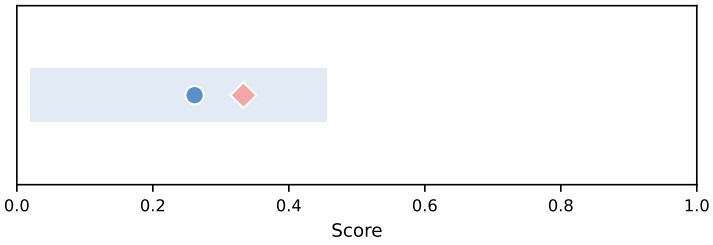


Caption: A showgirl puts on her lipstick in the mirror .
BEiTScore^{++} Median=0.038 **BEiTScore^{++} CI@90%**=[0.000, 0.083]
Human Rating=0.000



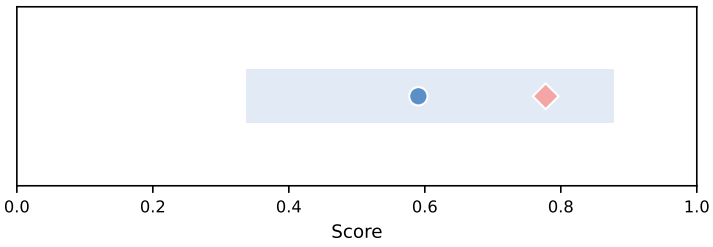


Caption: A sunset sky over rippling water .
BEiTScore^{++} Median=0.261 **BEiTScore^{++} CI@90%**=[0.019, 0.456]
Human Rating=0.333



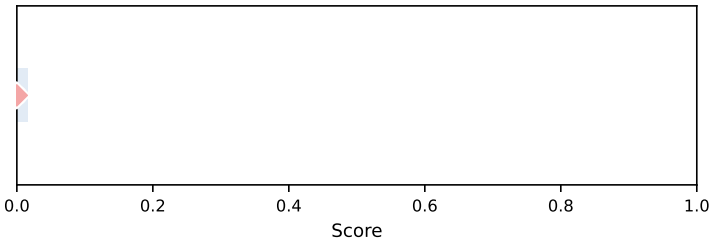


Caption: Person snowboarding at a ski slope .
BEiTScore^{++} Median=0.591 **BEiTScore^{++} CI@90%**=[0.338, 0.878]
Human Rating=0.778



Caption: Two dogs , one brown and white and one black and white , run on a field .
BEiTScore^{++} Median=0.000 **BEiTScore^{++} CI@90%**=[0.000, 0.016]
Human Rating=0.000





Caption: A person is hiking to the top of a hill .
BEiTScore^{++} Median=0.250 **BEiTScore^{++} CI@90%**=[0.013, 0.438]
Human Rating=0.111



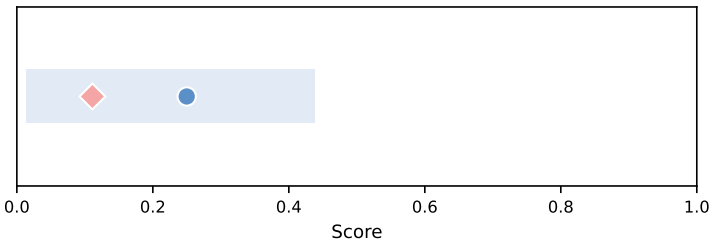
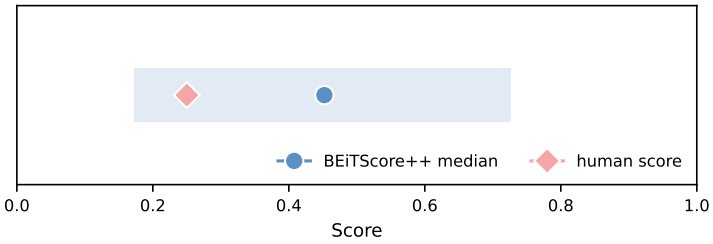


Figure 24: Additional qualitative examples for the predictions of BEiTScore++-L on the Ex-8K dataset.

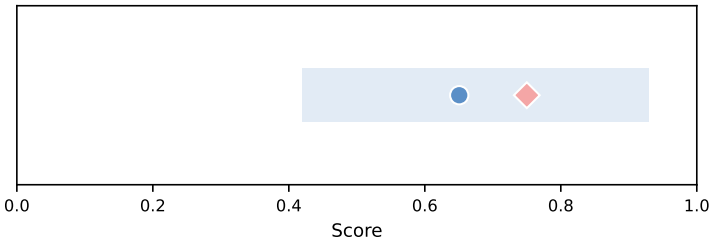


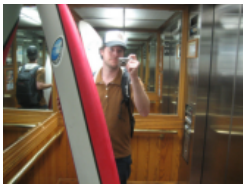
Caption: a cow is standing in the middle of a street .
BEiTScore^{++} Median=0.452 **BEiTScore^{++} CI@90%**=[0.172, 0.726]
Human Rating=0.250



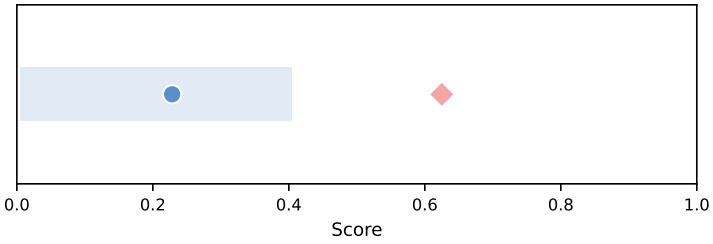


Caption: a man in a uniform costume stands next to a man in a horro costume.
BEiTScore^{++} Median=0.651 **BEiTScore^{++} CI@90%**=[0.420, 0.929]
Human Rating=0.750



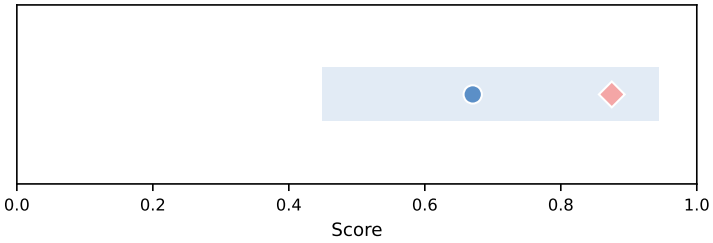


Caption: person is smiling.person is hanging train in the off.The scene contains train.
BEiTScore^{++} Median=0.228 **BEiTScore^{++} CI@90%**=[0.004, 0.404]
Human Rating=0.625





Caption: man in a coat walking with other people in the background.
BEiTScore^{++} Median=0.670 **BEiTScore^{++} CI@90%**=[0.449, 0.944]
Human Rating=0.875





Caption: a woman sitting on a bench next to a bicycle .
BEiTScore^{++} Median=0.047 **BEiTScore^{++} CI@90%**=[0.000, 0.098]
Human Rating=0.000

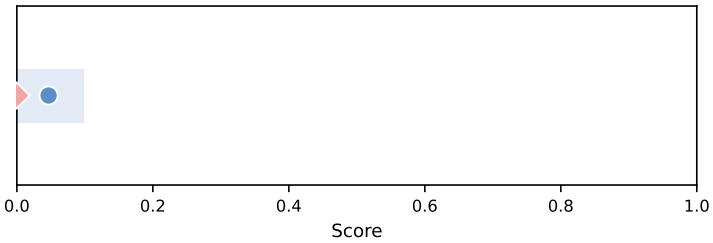
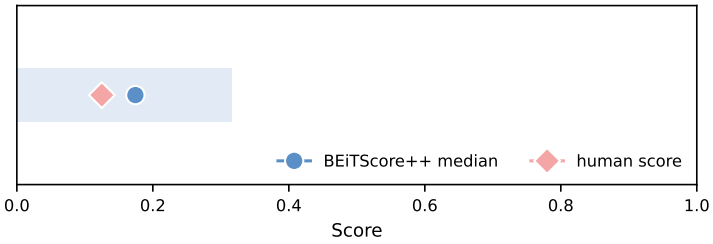


Figure 25: Qualitative examples for the predictions of BEiTScore++-L on the COM dataset. Each row shows the input image, the candidate caption, the predicted median and 90 percent interval, and the mean human rating.

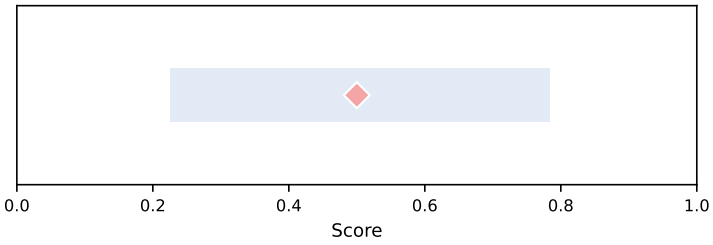


Caption: The scene contains displayed cakes.
BEiTScore^{++} Median=0.174 **BEiTScore^{++} CI@90%**=[0.000, 0.316]
Human Rating=0.125



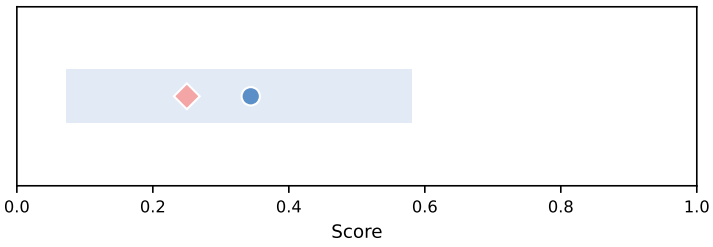


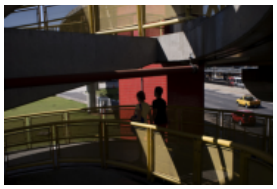
Caption: a group of people are standing in front of a building .
BEiTScore^{++} Median=0.500 **BEiTScore^{++} CI@90%**=[0.225, 0.783]
Human Rating=0.500



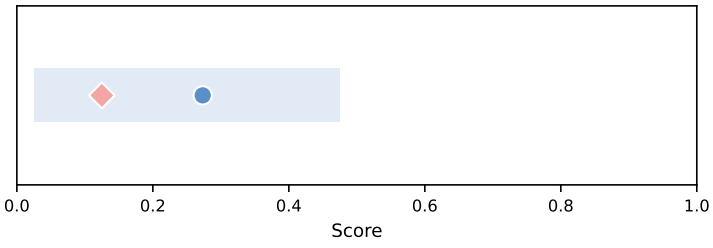


Caption: The scene contains booths.
BEiTScore^{++} Median=0.344 **BEiTScore^{++} CI@90%**=[0.072, 0.580]
Human Rating=0.250





Caption: a man in a white shirt is standing on a ladder .
BEiTScore^{++} Median=0.273 **BEiTScore^{++} CI@90%**=[0.025, 0.475]
Human Rating=0.125





Caption: dog is running.person is skateboarding.person is jumping.
BEiTScore^{++} Median=0.204 **BEiTScore^{++} CI@90%**=[0.000, 0.365]
Human Rating=0.250

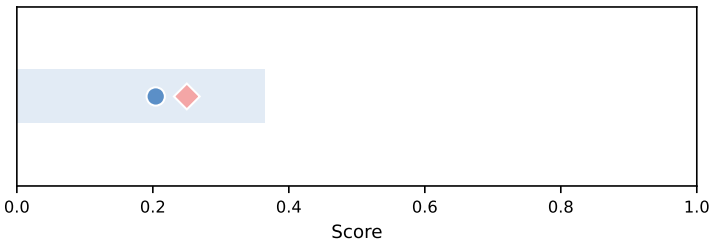


Figure 26: Additional qualitative examples for the predictions of BEiTScore++-L on the COM dataset.