

F1 is Not Enough! Models and Evaluation Towards User-Centered Explainable Question Answering

Hendrik Schuff^{1,2}

Heike Adel¹

Ngoc Thang Vu²

¹Bosch Center for Artificial Intelligence, Renningen, Germany

²Institut für Maschinelle Sprachverarbeitung, University of Stuttgart

{Hendrik.Schuff, Heike.Adel}@de.bosch.com

Thang.Vu@ims.uni-stuttgart.de

Abstract

Explainable question answering systems predict an answer together with an explanation showing why the answer has been selected. The goal is to enable users to assess the correctness of the system and understand its reasoning process. However, we show that current models and evaluation settings have shortcomings regarding the coupling of answer and explanation which might cause serious issues in user experience. As a remedy, we propose a hierarchical model and a new regularization term to strengthen the answer-explanation coupling as well as two evaluation scores to quantify the coupling. We conduct experiments on the HOTPOTQA benchmark data set and perform a user study. The user study shows that our models increase the ability of the users to judge the correctness of the system and that scores like F_1 are not enough to estimate the usefulness of a model in a practical setting with human users. Our scores are better aligned with user experience, making them promising candidates for model selection.

1 Introduction

Understanding the decisions of deep learning models is of utmost importance, especially when they are deployed in critical domains, such as medicine or finance (Ribeiro et al., 2016). In natural language processing (NLP), a variety of tasks have been addressed regarding explainability of neural networks, such as textual entailment (Camburu et al., 2018), sentiment classification (Clos et al., 2017), machine translation (Stahlberg et al., 2018) and question answering (Yang et al., 2018). In this paper, we address question answering (QA) due to its proximity to users in real-life settings, for instance, in the context of personal assistants.

Explainable question answering (XQA) is the task of (i) answering a question and (ii) providing an explanation that enables the user to understand

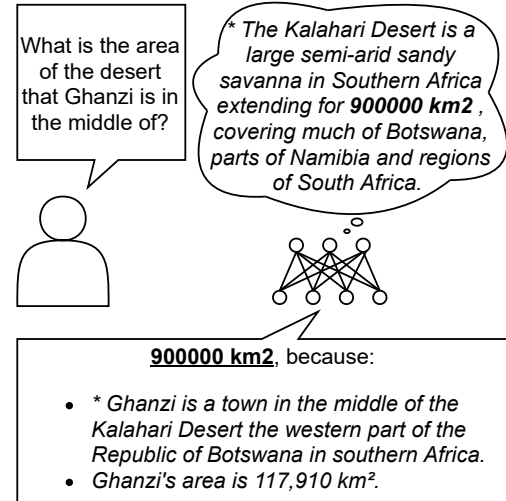


Figure 1: Example output of a representative XQA system (Yang et al., 2018) that would receive an answer- F_1 of 1 and an explanation- F_1 of 0.5 although the explanation provides no value to the user since the actual answer evidence (shown in cloud) is not included in the explanation (asterisks mark ground truth explanation).

why the answer was selected, e.g., by pointing to the facts that are needed for answering the question. Compared to approaches that output importance weights or analyze gradients (Simonyan et al., 2014; Ribeiro et al., 2016; Lundberg and Lee, 2017; Sundararajan et al., 2017), this has the advantage that the explanations are intuitively assessable even by lay users without machine learning background.

A *good explanation* (i.e., one that is helpful for the user) should therefore satisfy the following requirements: (i) It should contain all information that the model used to predict the answer for the question. This is necessary so that the user can reconstruct the model’s reasoning process. (ii) It should not include additional information that it did not use for predicting the answer. Otherwise, the explanation will confuse the users rather than help them. Note that these requirements do not

only hold for correct model decisions but are also valid for explaining wrong model answers so that the user can assess the correctness of the answers. Previous work on XQA mostly focuses on developing models that predict the correct answer and, independent of this, the correct explanation (Yang et al., 2018; Qi et al., 2019; Shao et al., 2020). This can lead to model outputs in which the explanations do not sufficiently relate to the answers. Consider the example provided in Figure 1. The model gives the correct answer to the question and provides an explanation consisting of one out of two relevant facts. However, the most important relevant fact (in which the answer actually appears) is not part of the explanation. As a result, the user cannot assess whether the model answer is correct or not and, thus, cannot trust the system. To strengthen the *coupling of answer and explanation prediction* in the model architecture and during training, we propose two novel approaches in this paper: (i) a hierarchical neural network architecture for XQA that ensures that only information included in the explanation is used to predict the answer to the question, and (ii) a regularization term for the loss function that explicitly couples answer and explanation prediction during training.

A good evaluation measure should score explanations by satisfying the following requirements: (i) It should reward explanations that are coupled to the answers of the model. (ii) It should punish explanations that are unrelated to the answers of the model. (iii) It should be correlated to user experience. Since explanations cannot only empower the user to assess the correctness of a system (Biran and McKeown, 2017; Kim et al., 2016) but also improve user satisfaction and confidence (Sinha and Swearingen, 2002; Biran and McKeown, 2017) and, thus, increase the acceptance of automatic systems (Herlocker et al., 2000; Cramer et al., 2008), this aspect is very important when evaluating models that should be applied in real-life scenarios. In most recent works, evaluation of XQA models focuses on optimizing F_1 -scores of answers and explanations (a collection of so-called supporting or relevant facts) (Yang et al., 2018). However, F_1 -scores only assess model outputs with respect to ground-truth annotations which only contain explanations for the correct answer. Thus, they fail to quantify the coherence between answer and explanation, especially when the predicted model answer is wrong. The example model output in

Figure 1 leads to an answer- F_1 -score of 1 and an explanation- F_1 -score of 0.5 although the explanation is useless for the user as described before. To *quantify the model’s answer-explanation coupling*, we propose two novel evaluation scores: (i) FARM which tracks prediction changes when removing facts, and (ii) LOCA which assesses whether the answer is contained in the explanation or not. Both scores do not require ground-truth annotations.

To summarize, we make contributions in two directions in this paper: For *modeling*, (i) we propose a hierarchical neural network architecture as well as (ii) a regularization term for the loss function of XQA systems. For *evaluation*, (iii) we propose two scores that are able to quantify a model’s answer-explanation coupling without relying on ground-truth annotations. (iv) To investigate the relation between different evaluation scores and user experience, we conduct a *user study*. The results show that our proposed models increase the ability of the user to judge the correctness of an answer and that our scores are stronger predictors of human behavior than standard scores like F_1 . (v) For reproducibility and future research, we will release code for our methods and for computing the evaluation scores as well as the user study data.¹

2 Related Work

In the context of XQA, Yang et al. (2018) present the HOTPOTQA data set which we also use for the experiments in this paper. In addition to questions and answers, it contains explanations in the form of relevant sentences from Wikipedia articles.

Most state-of-the-art models for HOTPOTQA extend the BiDaf++ architecture (Clark and Gardner, 2018; Seo et al., 2017), e.g., (Yang et al., 2018; Qi et al., 2019; Nishida et al., 2019; Ye et al., 2019; Qiu et al., 2019; Shao et al., 2020). Other approaches are based on question decomposition (Min et al., 2019; Perez et al., 2020), graph/hierarchical structures (Tu et al., 2019; Fang et al., 2019; Asai et al., 2020), virtual knowledge bases (Dhingra et al., 2020) or transformer models (Zhao et al., 2020).

So far, all of the research work on HOTPOTQA focuses on reaching higher F_1 -scores. In contrast, we question whether this actually aligns with user experience. To the best of our knowledge, only Chen et al. (2019) additionally conduct a human

¹<https://github.com/boschresearch/fl-is-not-enough>

evaluation. This confirms the observation of Adadi and Berrada (2018) that only very few papers related to explainable AI address (human) evaluation of explainability. Despite the large body of research in the field of human computer interaction, Abdul et al. (2018) show that there is a lack of collaboration and transfer of results to machine learning communities.

Another line of research our work relates to is the criticism of automatic evaluation scores. One frequently questioned score is BLEU (Papineni et al., 2002), which was shown to only correlate weakly with human judgements in tasks like machine translation (Callison-Burch et al., 2006), storytelling (Wang et al., 2018) and dialogue response generation (Liu et al., 2016). F_1 has been criticized from various perspectives including theoretical considerations and concrete applications (Hand and Christen, 2018; Chicco and Jurman, 2020; Sokolova et al., 2006). Qian et al. (2016) show that modifying F_1 -scores based on insights from psychometrics improves their correlation with human evaluations. In this paper we criticize the usage of F_1 as a measure of explainability in XQA and show in a user study that it is not related to user experience.

3 Methods for XQA

We built upon the model by Qi et al. (2019) as it is an improved version of the BiDaf++ model, which is used in numerous state-of-the-art XQA models (Yang et al., 2018; Qi et al., 2019; Nishida et al., 2019; Ye et al., 2019; Qiu et al., 2019) including the best-scoring publication (Shao et al., 2020). It consists of a question and context encoding part with self-attention, followed by two prediction heads: a prediction of relevant facts (i.e., the explanation) and a prediction of the answer to the question. The two heads are trained in a multi-task fashion based on the sum of their respective losses. First, we analyze the outputs of the model, revealing severe weaknesses in answer-explanation coupling. To address those weaknesses, we then propose (i) a novel neural network architecture that *selects* and *forgets* facts, and (ii) a novel answer-explanation coupling regularization term for the loss function.

3.1 Limitations of Current Models

We manually analyze outputs of the models by Qi et al. (2019) and Yang et al. (2018) and identify the following two problems.

Silent Facts. The models make use of facts without including them into their explanations (cf., Figure 1). As a result, the predicted answer does not occur in the explanation, leaving the user uninformed about where it came from.

Unused Facts. The models predict facts to be relevant without any relation to the predicted answer. The second fact of the explanation in Figure 1 is an example for this. We also found examples where the facts predicted to be relevant do not even contain the entities from the question.

3.2 Select & Forget Architecture

To explicitly ensure that the model only uses information from facts it predicts to be relevant for the answer selection, we propose a hierarchical model that first *selects* facts which are relevant to answer the question and then *forgets* about all other facts (see Figure 2). We use recurrent and self-attention layers to create encodings of the question and the context. In particular, we create two different encodings: one that will be used for predicting the relevance of the facts (fact-specific encoding) and one that will be used for predicting the answer to the question (QA-specific encoding). Based on the fact-specific encoding, the model first predicts which facts are relevant to answer the question. Next, we reduce the QA-specific encoding of the context based on the relevance predictions. In particular, we mask all facts that were not predicted to be relevant by zeroing out their encodings. The reduced context representation is concatenated with the QA-specific question encoding and passed on to the answer prediction, which we implement in the same way as Qi et al. (2019). Thus, the answer prediction now only receives encodings of facts that the model has predicted to be relevant. It predicts the type of the answer (yes/no/text span), as well as the start and end positions of the answer span within the context.

3.3 Answer-Fact Coupling Regularizer

Our second method addresses the coupling of answer and explanation prediction by modifying the loss function. The loss function used by Qi et al. (2019) is the sum of four cross entropy losses concerning (i) the answer type (yes/no/span) distribution, (ii) the answer start token distribution, (iii) the answer end token distribution, and (iv) the fact relevance distributions. All terms are optimized to be close to their respective ground truth annota-

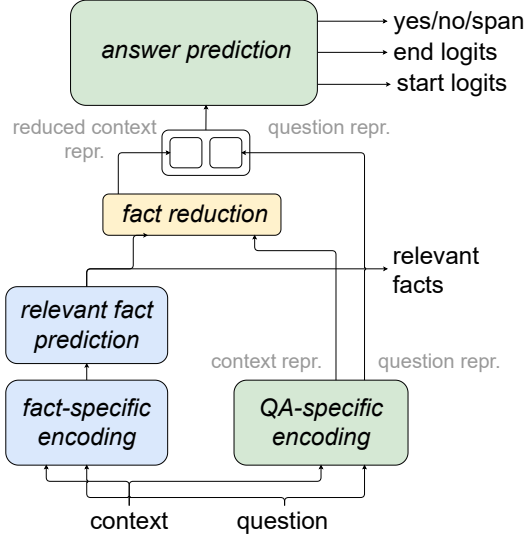


Figure 2: “Select and Forget” architecture with task-specific encodings and fact reduction.

tions. This is the desired effect in many, but, as our examples in Section 3.1 show, not in all situations. The loss function especially encourages the model to predict the ground truth explanation rather than an explanation that explains the predicted answer.

In order to reward a coupling between answer and explanation, we propose to add the following regularization term to the loss function:

$$J_{\text{reg}} = p_a \cdot \underbrace{\left(\overbrace{p_e \cdot 0}^{\text{GT expl.}} + \overbrace{(1 - p_e) \cdot c_1}^{\text{non-GT expl.}} \right)}_{\text{correct answer}} + \underbrace{\left(\overbrace{(1 - p_a) \cdot p_e \cdot c_2}^{\text{GT expl.}} + \overbrace{(1 - p_a) \cdot (1 - p_e) \cdot c_3}^{\text{non-GT expl.}} \right)}_{\text{wrong answer}} \quad (1)$$

with p_a corresponding to the probability of the model for the correct answer span and p_e denoting the probability of the model for the ground truth relevant facts. The term can be broken down into four cases: (i) correct answer and ground truth explanation, (ii) correct answer but non-ground truth explanation, (iii) incorrect answer but ground truth explanation and (iv) incorrect answer and non-ground truth explanation. Each case corresponds to a constant cost of 0, c_1 , c_2 and c_3 , respectively, with c_1, c_2, c_3 being hyperparameters. The resulting cost J_{reg} is the sum of the four individual costs weighted with their respective probabilities.

In particular, p_a is defined as the product of the probabilities assigned to start and end token positions of the answer span. For a data set instance with a context containing N facts, we define

$s_t \in \{0, 1\}^N$ as the ground truth annotations for the relevant facts. Accordingly, we denote the model’s relevance probability estimates with $s_p \in [0, 1]^N$. Based on this, we define p_e as $p_e^* = \prod_{i \in F} s_{p_i}$ with $F = \{i \in \{1, \dots, N\} : s_{t_i} = 1\}$ denoting the indices that correspond to ground truth facts. This corresponds to the joint probability of selecting the ground truth facts assuming the single selection probabilities to be independent. For our experiments, we adapt this definition to a numerically more stable term p_e^+ by replacing the product with a sum as this led to slightly better results on the development set.

4 New Evaluation Scores for XQA

In this section, we motivate that standard scores like F_1 are not enough to score XQA systems by presenting their limitations in those settings. To be able to quantify the degree to which a model is affected by those limitations, we propose two scores that go beyond standard scores: the *fact-removal score* and the *answer-location score*. Both scores can be calculated without any assumptions on model architecture and no need for ground truth annotations for answers or supporting facts.

4.1 Limitations of Current Evaluation Scores

Current evaluation of XQA is focused on three scores: (i) answer- F_1 , which is based on the token overlap between the predicted and the ground truth answer, (ii) SP- F_1 , which calculates F_1 based on the overlap of predicted and ground truth relevant (“supporting”) facts and (iii) joint- F_1 , which is based on the definitions of joint precision and joint recall as the products of answer and SP precision and recall as described in Yang et al. (2018). For HOTPOTQA, models are ranked based on joint- F_1 . We argue that this creates a false incentive that potentially hinders the development of truly usable models for the following reasons.

No Empirical Evidence. There is no empirical evidence that joint- F_1 is related to user performance or experience regarding XQA.

Rewarding Poor Explanations. Figure 1 shows an example prediction that is rewarded with a joint- F_1 of 0.5 although its explanation provides no value to the user. The reward stems from the overlap of the explanation with the ground truth but does not consider that the predicted answer is not contained in any of the predicted relevant facts.

Punishing Good Explanations. Consider a model output in which the predicted answer is wrong but the explanation perfectly explains this wrong answer, showing to the user why the model has selected it. Standard F_1 -scores compare the model output to the ground truth annotations and will, therefore, score both the answer and the explanation with an F_1 of 0. However, we argue that an explanation should be evaluated with a score higher than 0 if it is able to explain the reasoning process of the model to the user and, thus, lets the user identify the failure of the model.

4.2 Fact-Removal Score (FARM)

Ideally, the explanations of the model include all facts that the model uses within its reasoning chain but no additional facts beyond that. Note that even for a wrong model answer, this assumption should hold so that the relevant facts provide explanations for the (wrongly) predicted answer. To quantify the degree of answer-explanation coupling, we propose to iteratively remove parts (individual facts) of the explanation, re-evaluate the model using the reduced context and track how many of the model’s answers change. For a model with perfect coupling of answer and explanation, the answer will change with the first fact being removed (assuming no redundancy) but will not change when removing irrelevant facts not belonging to the explanation. We remove facts in order of decreasing predicted relevance as more relevant facts should influence the model’s reasoning process the strongest.

In the following, we denote an instance of the data set by $e \in E$ with its corresponding question e_{ques} and context e_{con} . We use $\text{answer}(\cdot, \cdot)$ to denote the answer that a model predicts for a given question and context. The functions $\text{reduce}_{\text{rel}}(\cdot, k)$ ($\text{reduce}_{\text{irr}}(\cdot, k)$) return a context from which up to k facts the model predicts to be relevant (irrelevant) have been removed.² We re-evaluate the model on this reduced context and calculate the fraction of changed answers $c_{\text{rel}}(k)$ and $c_{\text{irr}}(k)$, respectively.

$$a(e) = \text{answer}(e_{\text{ques}}, e_{\text{con}}) \quad (2)$$

$$\hat{a}_{\text{rel},k}(e) = \text{answer}(e_{\text{ques}}, \text{reduce}_{\text{rel}}(e_{\text{con}}, k)) \quad (3)$$

$$\hat{a}_{\text{irr},k}(e) = \text{answer}(e_{\text{ques}}, \text{reduce}_{\text{irr}}(e_{\text{con}}, k)) \quad (4)$$

$$c_{\text{rel}}(k) = \frac{|\{e \in E : a(e) \neq \hat{a}_{\text{rel},k}(e)\}|}{|E|} \quad (5)$$

$$c_{\text{irr}}(k) = \frac{|\{e \in E : a(e) \neq \hat{a}_{\text{irr},k}(e)\}|}{|E|} \quad (6)$$

²If the number of facts predicted as (ir)relevant is less or equal to k , we remove all (ir)relevant facts from the context.

Finally, we condense $c_{\text{rel}}(k)$ and $c_{\text{irr}}(k)$ into a single fact-removal score:

$$\text{FARM}(k) = \frac{c_{\text{rel}}(k)}{1 + c_{\text{irr}}(k)} \in [0, 1] \quad (7)$$

$\text{FARM}(k)$ ranges between zero and one and a higher score corresponds to a better explanation.

4.3 Answer-Location Score (LOCA)

A second important indicator for the degree of a model’s answer-explanation coupling is the location of the answer span: As shown in Figure 1, the models can predict answers that are located outside the facts they predict to be relevant, i.e., outside the explanation. This is confusing for a user. Therefore, we consider the fractions of answer spans that are inside the explanation of the model and the fraction of answer spans that are outside. For an ideal model, all answer spans would be located inside the explanation. We use I and O to denote the number of answers inside/outside of the set of facts predicted as relevant. A denotes the total number of answers.³ Based on these counts, we propose the answer-location score that we define as

$$\text{LOCA} = \frac{\frac{I}{A}}{1 + \frac{O}{A}} = \frac{I}{A + O} \in [0, 1]. \quad (8)$$

The LOCA score ranges between zero and one, with larger values indicating better answer-explanation coupling.

5 Experiments and Results

In this section, we describe the dataset we used in our experiments as well as our results. More details for reproducibility, including hyperparameters, are provided in the appendix.

5.1 Dataset

The HOTPOTQA data set is a multi-hop open-domain explainable question answering data set containing 113k questions with crowd-sourced annotations. Each instance of the training data contains a question, a context consisting of the first paragraph of ten Wikipedia articles, the annotated answer and an explanation in the form of a selection of relevant sentences from the context. As HOTPOTQA was designed as a multi-hop data set,

³In HOTPOTQA, answers can stem from article titles although titles are never used as relevant facts. Thus, $A > I + O$ is possible. Our score is still applicable in this case.

	Score	Qi-2019	S&F	reg.	S&F+reg.
Standard Scores	Answer-EM	49.48	46.09	49.67	46.44
	Answer-F ₁	63.76	59.99	63.56	60.60
	Answer-P	66.26	62.41	66.27	62.77
	Answer-R	65.52	61.61	65.06	62.53
	SP-EM	39.81	42.16	25.98	30.45
	SP-F ₁	79.34	80.07	75.60	77.53
	SP-P	78.01	78.84	66.79	70.09
	SP-R	85.26	85.45	93.26	92.03
	Joint-EM	22.28	22.78	14.56	16.38
	Joint-F ₁	52.51	50.71	49.66	48.99
	Joint-P	53.33	51.40	45.64	45.74
	Joint-R	57.92	55.61	62.09	59.52
Proposed Scores	FARM(4)	66.20	75.54	73.32	76.64
	↳ c _{rel} (4)	77.06	86.05	81.69	84.58
	↳ c _{irrel} (4)	16.39	13.91	11.41	10.36
	LOCA	60.49	70.60	67.92	75.56
	↳ I	67.48	71.68	72.60	76.32
	↳ O	11.55	1.53	6.89	1.01

Table 1: Comparison of our methods to Qi et al. (2019) regarding evaluation scores from related work and our proposed scores on the distractor dev set (SP: supporting facts). All values in %.

finding the answer to a question requires combining information from two different articles. The eight other articles are *distracting* the system.⁴

5.2 Experimental Results

In our experiments, we assess the effects of our Select & Forget architecture (S&F) and the regularization term (reg.). Table 1 shows our approaches in comparison to the model by Qi et al. (2019).⁵ While our S&F architecture performs comparable in standard scores like answer-exact-match (Answer-EM), answer-F₁, joint-EM and joint-F₁ (for some of them slightly better, for some of them slightly lower), the regularization term increases the recall of the relevant fact prediction considerably.

In terms of our proposed scores for measuring answer-explanation-coupling, all our three models clearly outperform the baseline model (lower part of Table 1). In the first three rows, we report the models’ FARM scores and the fractions of changed answers for $k = 4$, i.e., when a maximum of four facts are removed. We choose $k = 4$ as this is the

⁴The data set also contains a *full wiki* test set in which the context spans all collected Wikipedia articles. We focus on the distractor setting in this paper. The data set can be downloaded from <https://hotpotqa.github.io/>.

⁵We retrain their model using the implementation and preprocessing provided at <https://github.com/qipeng/golden-retriever>.

highest number of facts within an explanation in the ground truth annotations of the HOTPOTQA data. The last three rows show the LOCA scores and the respective fractions of answers inside and outside facts predicted as relevant.

The different behavior of models regarding joint-F₁ vs. FARM and LOCA raises the question which scores are better suited to quantify explainability in a real-life setting with human users. To answer this, we conduct a user study in Section 6.

6 Human Evaluation

We conduct a user study to investigate whether standard scores like F₁ or our proposed scores are better suited to predict user behavior and performance. Moreover, the study provides another way to compare our proposed methods to the model by Qi et al. (2019) and the ground truth explanations. In contrast to the human evaluation from Chen et al. (2019), we evaluate explanations in the context of the model answer, ask participants to rate the predictions along multiple dimensions and collect responses from 40 instead of 3 subjects.

6.1 Choice of Models

We choose to compare the model proposed by Qi et al. (2019) (called “Qi-2019” in the following) as a representative of the commonly used BiDaf++ architecture (Clark and Gardner, 2018) in XQA, our proposed Select & Forget architecture (S&F) and our proposed regularization term (reg.). In addition, we include the ground truth (GT) annotations to set an upper bound. Although the combination of regularization and the S&F architecture reaches promising performance in Table 1, we assess the effects of our methods in isolation here and leave the evaluation of the combination to future work.

6.2 Study Design

We make use of a unifactorial between-subject design in which each model constitutes one condition. We randomly sample a set of 25 questions from the HOTPOTQA dev set and collect the model answer and explanation predictions (or annotations for GT) for each condition. For each answer prediction, we manually assess whether it is equivalent to the ground truth answer.⁶ Each participant sees the 25 questions and the answers and explanations

⁶For example, we consider the answer prediction “Firth of Clyde, Scotland” to be equivalent to the ground truth answer “Firth of Clyde”.

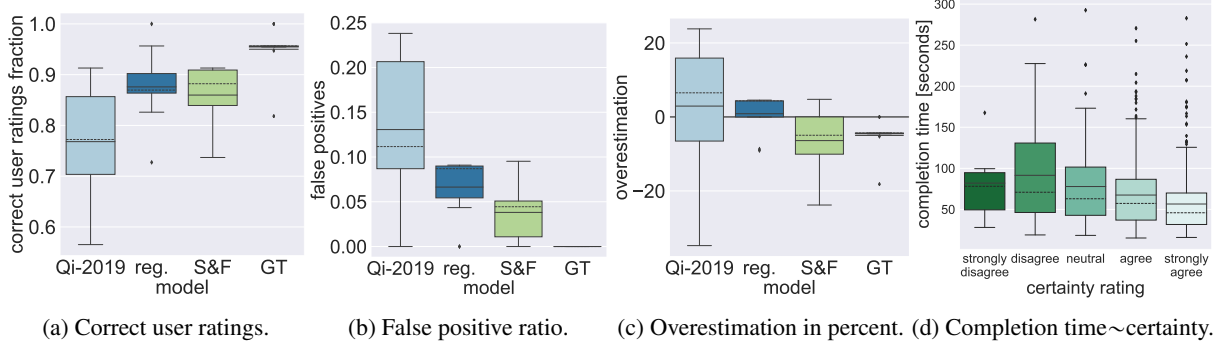


Figure 3: Boxplots showing results from the user study. Boxes mark quartiles, whiskers mark 1.5 inter-quartile ranges, outliers are plotted separately. Horizontal solid/dashed lines within boxes mark means and medians, resp.

of one model in a random permutation. For each question, we ask the participants to rate whether the model answer is correct. In addition, we ask for multiple self-reports to assess, e.g., the trust of the user in the system. In particular, we track the variables discussed in the following subsection.

6.3 Dependent Variables

We derive multiple dependent variables from the participants’ ratings, namely completion time (Lim et al., 2009; Lage et al., 2019), several performance variables indicating how well they judged the correctness of the model (fraction of correct ratings, false positive ratio (FP), false negative ratio (FN), true positive ratio (TP), true negative ratio (TN), precision (P), recall (R) and F_1 values), agreement (fraction of model predictions that the users rate as correct (Bussone et al., 2015)), and overestimation (difference between agreement and true model accuracy (Nourani et al., 2019)).

Furthermore, we collect the following variables in self-reports with five-point Likert scales: certainty of the participants (Greis et al., 2017a), completeness and helpfulness of the explanations (Nourani et al., 2019), trust of the participants in the model (Bussone et al., 2015), and satisfaction (Kulesza et al., 2012; Greis et al., 2017b).

All the questions and screenshots of our study are given in the appendix.

6.4 Participants and Data Cleaning

We collect the ratings of 40 participants (16 female, 24 male) with a mean age of 26.6 years ($SD = 3.4$). We filter out all responses with a completion time smaller than 15 seconds or larger than 5 minutes as this indicates that the participant did not read the whole explanation or was interrupted during the study. We further asked them

whether they knew the answer before and exclude the responses to those questions from our evaluation. In total, we discard 12.10% of the responses.

6.5 Results

In this section, we summarize the main results of the user study. For better overview, we do not include evaluations on every variable from Section 6.3 but show them in the appendix.

Figure 3a shows the fraction of correct user ratings of model correctness. The correctness of our proposed models can be better judged than the Qi-2019 model: The regularized model and the S&F model increase the fraction by 10.79% and 9.17%, respectively, compared to Qi-2019. The GT comparison shows an upper bound.

Among the performance variables, the false positive ratio deserves particular attention as a false positive corresponds to a user thinking the model answer is correct while it is not. Such an error can be dangerous in safety-critical domains. Figure 3b shows that the fraction of FPs is decreased by 6.43% by the regularized model and by 9.25% by the S&F model compared to Qi-2019. The ground truth has zero false positives by definition.

A similar effect can be seen when evaluating overestimation: Both our models alleviate overestimation as shown in Figure 3c. While participants overestimate the model accuracy of Qi-2019 by 2.93% on average, the regularized model only leads to 0.87% overestimation. The S&F model is even underestimated by 6.40% on average. While an ideal model would lead to neither over- nor underestimation, underestimation can be preferable to overestimation if the model is deployed into high-risk environments, such as medical contexts. In general, a reduction in overestimation can — besides enhanced fact selection — also be linked to

Human eval.	Standard Scores												Proposed Scores						
	Answer				Supporting Facts				Joint				Answer Changes		FARM		%in-fact		LOCA
	EM	F ₁	P	R	EM	F ₁	P	R	EM	F ₁	P	R	rel.	irrel.	FARM(1)	FARM(4)	rel.	irrel.	LOCA
correct decision																			
overestimation		+		+															
completion time	+		+																
human-FP		+		+															
human-TP																			
human-FN																			
human-TN																			
human-P																			
human-R																			
human-F ₁																			

Table 2: The table shows whether sorting the conditions by a human score (rows) and an automatized score (columns) results in the same order (+), the inverse order (-) or a different order (blank cell). Green (■) cells with circles mark desirable relations, red (■) cells without circles mark undesirable relations.

an improved answer accuracy as better-performing models naturally leave more room for underestimation.

Finally, we consider relations within the variables from Section 6.3. Figure 3d shows that mean completion time monotonously decreases with increasing user certainty (with the exception of “strongly disagree”).⁷ This confirms the findings of Greis et al. (2017a) who investigate the effect of user uncertainty on behavioral measurements.

6.6 Correlation with Evaluation Scores

Finally, we investigate the correlation of human ratings with model evaluation scores. We rank the models by (i) human measures obtained in the user study and (ii) model evaluation scores. In Table 2, a cell is marked with a “+” if the ranking with respect to the human measure and the model score is identical (e.g., the ranks regarding human-FPs and answer-F₁ are identical). If the ranks are exactly reversed, we mark the cell with a “-”. All other cells are left empty. “+” and “-” both indicate a perfect correlation and do not imply one being preferable over the other. Next, we consider whether selecting a model based on the different model scores would result in a desired change in human evaluation scores or not. This depends on whether a high score (e.g., F₁) or a low score (e.g., the fraction of answers outside the predicted relevant facts) is aimed for. We indicate desired model selection with green circled cells and undesired model selection with red cells (e.g., choosing a model with a higher answer-F₁ would result in a

⁷The low completion time for “strongly disagree” could indicate that the users could not find any relation at all between answer and explanation.

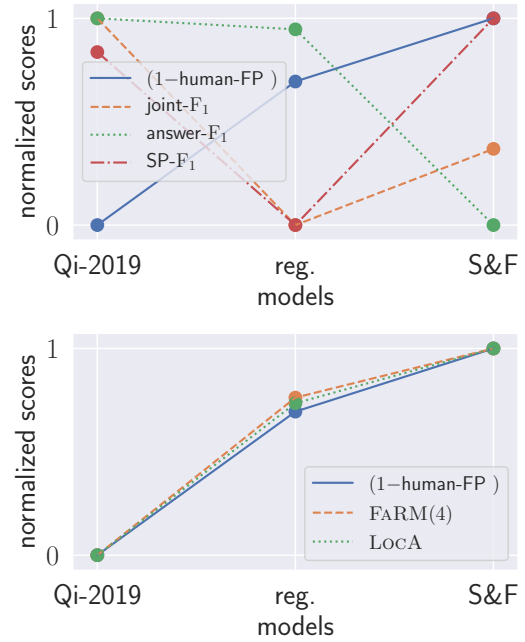


Figure 4: Model score comparisons between human-false positives and model scores. All scores are normalized to [0, 1]. We show (1-human-FP) as less FPs are better. The upper figure shows that F₁ poorly correlates to human performance. The lower figure shows a much stronger correlation for our proposed scores.

model with more human-FPs. This is not desired.) All F₁-scores show at least one undesirable rank relation. Notably, joint-F₁ is among the least aligned scores. In contrast, our scores have only desirable relations. In particular, FARM(4) and LOCA lead to a model ranking that is inverse to the ranking by human-overestimation and human-FPs. This is also confirmed in Figure 4, which shows how the human-FP ratio varies in comparison to the three F₁ scores (upper plot: no correlation) and to our pro-

posed scores (lower plot: correlated). See appendix for other dependent variables.

To sum up, our results indicate that (i) F_1 is not suited to quantify the explanatory power of a model and (ii) our proposed scores predict user behavior better than standard scores, opening the possibility of using them for model selection.

7 Conclusion

In this paper, we investigated explainable question answering, revealing that existing models lack an explicit coupling of answers and explanations and that evaluation scores used in related work fail to quantify that. This highly impairs their applicability in real-life scenarios with human users. As a remedy, we addressed both modeling and evaluation, proposing a hierarchical neural architecture, a regularization term, as well as two new evaluation scores. Our user study showed that our models help the users assess their correctness and that our proposed evaluation scores are better correlated with user experience than standard measures like F_1 .

Acknowledgement

We thank the members of the BCAI NLP&KRR research group and the anonymous reviewers for their helpful comments. Ngoc Thang Vu is funded by Carl Zeiss Foundation.

References

- Ashraf Abdul, Jo Vermeulen, Danding Wang, Brian Y. Lim, and Mohan Kankanhalli. 2018. [Trends and trajectories for explainable, accountable and intelligible systems: An hci research agenda](#). In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*, page 1–18, New York, NY, USA. Association for Computing Machinery.
- Amina Adadi and Mohammed Berrada. 2018. [Peeking inside the black-box: A survey on explainable artificial intelligence \(XAI\)](#). *IEEE Access*, 6:52138–52160.
- Akari Asai, Kazuma Hashimoto, Hannaneh Hajishirzi, Richard Socher, and Caiming Xiong. 2020. Learning to retrieve reasoning paths over wikipedia graph for question answering. In *8th International Conference on Learning Representations*, Addis Ababa, Ethiopia.
- Or Biran and Kathleen R. McKeown. 2017. [Human-centric justification of machine learning predictions](#). In *Proceedings of the Twenty-Sixth International Joint Conference on Artificial Intelligence*, pages 1461–1467, Melbourne, Australia. International Joint Conferences on Artificial Intelligence.
- Adrian Bussone, Simone Stumpf, and Dympha O’Sullivan. 2015. [The role of explanations on trust and reliance in clinical decision support systems](#). In *2015 International Conference on Healthcare Informatics*, pages 160–169, Dallas, TX, USA. IEEE Computer Society.
- Chris Callison-Burch, Miles Osborne, and Philipp Koehn. 2006. [Re-evaluating the role of Bleu in machine translation research](#). In *11th Conference of the European Chapter of the Association for Computational Linguistics*, Trento, Italy. Association for Computational Linguistics.
- Oana-Maria Camburu, Tim Rocktäschel, Thomas Lukasiewicz, and Phil Blunsom. 2018. [e-snli: Natural language inference with natural language explanations](#). In *Annual Conference on Neural Information Processing Systems 2018*, pages 9560–9572, Montréal, Canada. Advances in Neural Information Processing Systems.
- Jifan Chen, Shih-Ting Lin, and Greg Durrett. 2019. [Multi-hop question answering via reasoning chains](#). *Computing Research Repository*, arXiv:1910.02610.
- Davide Chicco and Giuseppe Jurman. 2020. [The advantages of the matthews correlation coefficient \(mcc\) over f1 score and accuracy in binary classification evaluation](#). *BMC genomics*, 21(1):6.
- Christopher Clark and Matt Gardner. 2018. [Simple and effective multi-paragraph reading comprehension](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 845–855, Melbourne, Australia. Association for Computational Linguistics.
- Jérémie Clos, Nirmalie Wiratunga, and Stewart Massie. 2017. Towards explainable text classification by jointly learning lexicon and modifier terms. In *IJCAI-17 Workshop on Explainable AI (XAI)*, page 19.
- Henriette S. M. Cramer, Vanessa Evers, Satyan Ramalal, Maarten van Someren, Lloyd Rutledge, Natalia Stash, Lora Aroyo, and Bob J. Wielinga. 2008. [The effects of transparency on trust in and acceptance of a content-based art recommender](#). *User Model. User-Adapt. Interact.*, 18(5):455–496.
- Bhuwan Dhingra, Manzil Zaheer, Vidhisha Balachandran, Graham Neubig, Ruslan Salakhutdinov, and William W. Cohen. 2020. Differentiable reasoning over a virtual knowledge base. In *8th International Conference on Learning Representations, ICLR 2020*, Addis Ababa, Ethiopia.
- Yuwei Fang, Siqi Sun, Zhe Gan, Rohit Pillai, Shuohang Wang, and Jingjing Liu. 2019. [Hierarchical graph network for multi-hop question answering](#). *Computing Research Repository*, arXiv:1911.03631.

- Miriam Greis, Jakob Karolus, Hendrik Schuff, Paweł W. Woźniak, and Niels Henze. 2017a. [Detecting uncertain input using physiological sensing and behavioral measurements](#). In *Proceedings of the 16th International Conference on Mobile and Ubiquitous Multimedia, MUM '17*, page 299–304, New York, NY, USA. Association for Computing Machinery.
- Miriam Greis, Hendrik Schuff, Marius Kleiner, Niels Henze, and Albrecht Schmidt. 2017b. [Input controls for entering uncertain data: Probability distribution sliders](#). volume 1, New York, NY, USA. Association for Computing Machinery.
- David Hand and Peter Christen. 2018. [A note on using the f-measure for evaluating record linkage algorithms](#). *Stat. Comput.*, 28(3):539–547.
- Jonathan L. Herlocker, Joseph A. Konstan, and John Riedl. 2000. [Explaining collaborative filtering recommendations](#). In *Proceedings of the 2000 ACM Conference on Computer Supported Cooperative Work, CSCW '00*, page 241–250, New York, NY, USA. Association for Computing Machinery.
- Been Kim, Oluwasanmi Koyejo, and Rajiv Khanna. 2016. [Examples are not enough, learn to criticize! Criticism for interpretability](#). In *Annual Conference on Neural Information Processing Systems 2016*, pages 2280–2288, Barcelona, Spain. Advances in Neural Information Processing Systems.
- Todd Kulesza, Simone Stumpf, Margaret Burnett, and Irwin Kwan. 2012. [Tell me more? The effects of mental model soundness on personalizing an intelligent agent](#). In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems, CHI '12*, page 1–10, New York, NY, USA. Association for Computing Machinery.
- Isaac Lage, Emily Chen, Jeffrey He, Menaka Narayanan, Been Kim, Samuel J Gershman, and Finale Doshi-Velez. 2019. [Human evaluation of models built for interpretability](#). In *Proceedings of the AAAI Conference on Human Computation and Crowdsourcing*, volume 7, pages 59–67.
- Brian Y. Lim, Anind K. Dey, and Daniel Avrahami. 2009. [Why and why not explanations improve the intelligibility of context-aware intelligent systems](#). In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems, CHI '09*, page 2119–2128, New York, NY, USA. Association for Computing Machinery.
- Chia-Wei Liu, Ryan Lowe, Iulian Serban, Mike Noseworthy, Laurent Charlin, and Joelle Pineau. 2016. [How NOT to evaluate your dialogue system: An empirical study of unsupervised evaluation metrics for dialogue response generation](#). In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 2122–2132, Austin, Texas. Association for Computational Linguistics.
- Scott M. Lundberg and Su-In Lee. 2017. [A unified approach to interpreting model predictions](#). In *Annual Conference on Neural Information Processing Systems 2017*, pages 4765–4774, Long Beach, CA, USA. Advances in Neural Information Processing Systems.
- Sewon Min, Victor Zhong, Luke Zettlemoyer, and Hananeh Hajishirzi. 2019. [Multi-hop reading comprehension through question decomposition and rescoring](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 6097–6109, Florence, Italy. Association for Computational Linguistics.
- Kosuke Nishida, Kyosuke Nishida, Masaaki Nagata, Atsushi Otsuka, Itsumi Saito, Hisako Asano, and Junji Tomita. 2019. [Answering while summarizing: Multi-task learning for multi-hop QA with evidence extraction](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 2335–2345, Florence, Italy. Association for Computational Linguistics.
- Mahsan Nourani, Samia Kabir, Sina Mohseni, and Eric D Ragan. 2019. [The effects of meaningful and meaningless explanations on trust and perceived system accuracy in intelligent systems](#). In *Proceedings of the AAAI Conference on Human Computation and Crowdsourcing*, volume 7, pages 97–105.
- Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. [Bleu: a method for automatic evaluation of machine translation](#). In *Proceedings of the 40th Annual Meeting of the Association for Computational Linguistics*, pages 311–318, Philadelphia, Pennsylvania, USA. Association for Computational Linguistics.
- Ethan Perez, Patrick Lewis, Wen-tau Yih, Kyunghyun Cho, and Douwe Kiela. 2020. [Unsupervised question decomposition for question answering](#). *Computing Research Repository*, arXiv:2002.09758.
- Peng Qi, Xiaowen Lin, Leo Mehr, Zijian Wang, and Christopher D. Manning. 2019. [Answering complex open-domain questions through iterative query generation](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 2590–2602, Hong Kong, China. Association for Computational Linguistics.
- Peng Qian, Xipeng Qiu, and Xuanjing Huang. 2016. [A new psychometric-inspired evaluation metric for Chinese word segmentation](#). In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2185–2194, Berlin, Germany. Association for Computational Linguistics.
- Lin Qiu, Yunxuan Xiao, Yanru Qu, Hao Zhou, Lei Li, Weinan Zhang, and Yong Yu. 2019. [Dynamically fused graph network for multi-hop reasoning](#).

- In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 6140–6150, Florence, Italy. Association for Computational Linguistics.
- Marco Tulio Ribeiro, Sameer Singh, and Carlos Guestrin. 2016. “Why should I trust you?”: Explaining the predictions of any classifier. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, KDD ’16, page 1135–1144, New York, NY, USA. Association for Computing Machinery.
- Min Joon Seo, Aniruddha Kembhavi, Ali Farhadi, and Hannaneh Hajishirzi. 2017. Bidirectional attention flow for machine comprehension. In *5th International Conference on Learning Representations, ICLR 2017*, Toulon, France.
- Nan Shao, Yiming Cui, Ting Liu, Shijin Wang, and Guoping Hu. 2020. Is graph structure necessary for multi-hop reasoning? *Computing Research Repository*, arXiv:2004.03096.
- Karen Simonyan, Andrea Vedaldi, and Andrew Zisserman. 2014. Deep inside convolutional networks: Visualising image classification models and saliency maps. In *2nd International Conference on Learning Representations, ICLR 2014*, Banff, AB, Canada.
- Rashmi Sinha and Kirsten Swearingen. 2002. The role of transparency in recommender systems. In *CHI ’02 Extended Abstracts on Human Factors in Computing Systems*, CHI EA ’02, page 830–831, New York, NY, USA. Association for Computing Machinery.
- Marina Sokolova, Nathalie Japkowicz, and Stan Szpakowicz. 2006. Beyond accuracy, f-score and ROC: A family of discriminant measures for performance evaluation. In *AI 2006: Advances in Artificial Intelligence, 19th Australian Joint Conference on Artificial Intelligence*, volume 4304 of *Lecture Notes in Computer Science*, pages 1015–1021, Hobart, Australia. Springer.
- Felix Stahlberg, Danielle Saunders, and Bill Byrne. 2018. An operation sequence model for explainable neural machine translation. In *Proceedings of the 2018 EMNLP Workshop BlackboxNLP: Analyzing and Interpreting Neural Networks for NLP*, pages 175–186, Brussels, Belgium. Association for Computational Linguistics.
- Mukund Sundararajan, Ankur Taly, and Qiqi Yan. 2017. Axiomatic attribution for deep networks. In *Proceedings of the 34th International Conference on Machine Learning, ICML 2017*, volume 70 of *Proceedings of Machine Learning Research*, pages 3319–3328, Sydney, NSW, Australia.
- Ming Tu, Kevin Huang, Guangtao Wang, Jing Huang, Xiaodong He, and Bowen Zhou. 2019. Select, answer and explain: Interpretable multi-hop reading comprehension over multiple documents. *Computing Research Repository*, arXiv:1911.00484.
- Xin Wang, Wenhui Chen, Yuan-Fang Wang, and William Yang Wang. 2018. No metrics are perfect: Adversarial reward learning for visual storytelling. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 899–909, Melbourne, Australia. Association for Computational Linguistics.
- Zhilin Yang, Peng Qi, Saizheng Zhang, Yoshua Bengio, William Cohen, Ruslan Salakhutdinov, and Christopher D. Manning. 2018. HotpotQA: A dataset for diverse, explainable multi-hop question answering. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 2369–2380, Brussels, Belgium. Association for Computational Linguistics.
- Deming Ye, Yankai Lin, Zhenghao Liu, Zhiyuan Liu, and Maosong Sun. 2019. Multi-paragraph reasoning with knowledge-enhanced graph neural network. *Computing Research Repository*, arXiv:1911.02170.
- Chen Zhao, Chenyan Xiong, Corby Rosset, Xia Song, Paul N. Bennett, and Saurabh Tiwary. 2020. Transformer-xh: Multi-evidence reasoning with extra hop attention. In *8th International Conference on Learning Representations, ICLR 2020*, Addis Ababa, Ethiopia.

Appendix

A HOTPOTQA Data Set

We conduct our analysis and experiments on the explainable multi-hop reasoning question answering data set HOTPOTQA (Yang et al., 2018). The data set contains 113k questions with crowd-sourced annotations for answers and explanations. Explanations correspond to a selection of supporting facts (i.e., sentences) from Wikipedia articles. The data set contains 90,564 training instances as well as 7405 development instances. The test set is split into two separate test sets with 7405 instances each: the full-wiki test set and the distractor test set. Models evaluated on the full-wiki test data need to retrieve relevant articles from a given set of Wikipedia articles (therefore *full wiki*) while the distractor test set provides models with ten articles of which two contain the supporting facts and the other eight articles are *distracting* the system. Both test sets are not publicly available. F₁ values on the test sets can be obtained by submitting a model to the leaderboard. As our proposed FARM and LOCA scores need to access context information, we evaluate all tested models on the distractor development set while training them on the provided training data. The training and development data as well as the leaderboard can be found online.⁸

⁸<https://hotpotqa.github.io/>

B S&F Workflow Example

Figure 5 shows an exemplary workflow of the S&F architecture. First, relevant facts are selected. Second, all other facts are masked and third, the answer is predicted based on the masked context.

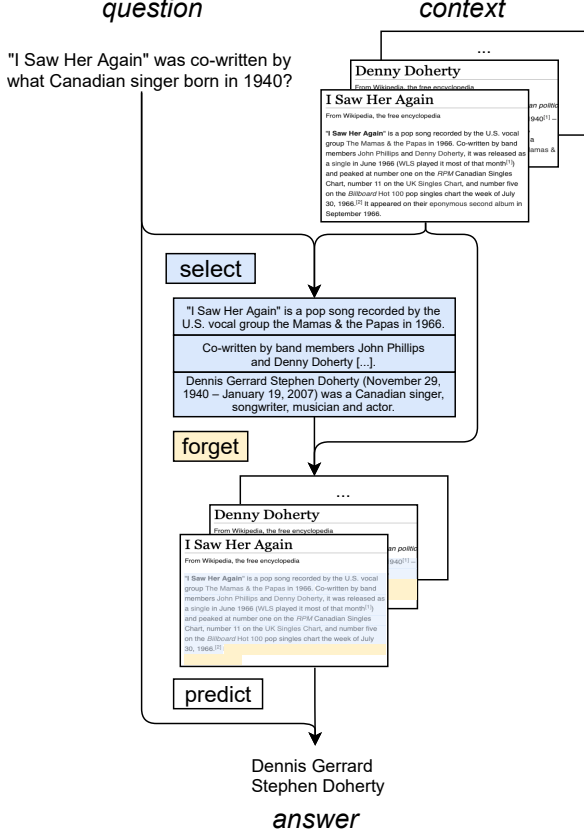


Figure 5: Exemplary workflow of the S&F model.

C Training Details and Hyperparameters

We use the same preprocessing, hyperparameters and early stopping procedure as Qi et al. (2019) to train our models.⁹ The Qi-2019 model as well as our regularized adaption contain 99M parameters, our S&F model and its regularized version contain 100M parameters each. The additional hyperparameters of our regularization term are optimized with random search using 100 runs. In particular, we sample $c_i \sim U([0.0, 5.0])$, $i \in \{1, 2, 3\}$ and select p_e^* and p_e^+ with equal probability. We select the best model based on the percentage of answer spans inside facts predicted as relevant to ensure decent answer-explanation coupling while not directly optimizing on the LOCA. Based on our

⁹<https://github.com/qipeng/golden-retriever>

	Metric	Yang-2018	Qi-2019
Standard Scores	Answer-EM	43.74	49.48
	Answer-F ₁	57.29	63.76
	Answer-P	59.76	66.26
	Answer-R	58.74	65.52
	SP-EM	24.54	39.81
	SP-F ₁	68.02	79.34
	SP-P	69.86	78.01
	SP-R	72.90	85.26
	Joint-EM	12.83	22.28
	Joint-F ₁	41.12	52.51
Proposed Scores	Joint-P	43.43	53.33
	Joint-R	45.44	57.92
	FARM(4)	50.08	66.20
	↳ $c_{rel}(4)$	62.50	77.06
	↳ $c_{irrel}(4)$	24.81	16.39
	LOCA	44.78	60.49
	↳ I	54.90	67.48
	↳ O	22.60	11.55

Table 3: Comparison of the HOPOTQA baseline model by Yang et al. (2018) and the modified model by Qi et al. (2019). The modified model outperforms the baseline on all scores. All values in %.

hyperparameter search, the best regularization parameters for the model proposed by Qi et al. (2019) are $c_1 = 4.96$, $c_2 = 2.02$ and $c_3 = 3.10$. The best parameters for the regularized S&F model are $c_1 = 1.18$, $c_2 = 0.24$ and $c_3 = 1.61$. We trained all models on Nvidia Tesla V100 GPUs.

D Comparison of BiDaf++ Versions

We compare the BiDaf++ model used as the official HOPOTQA baseline of Yang et al. (2018) to the modified model by Qi et al. (2019) in Table 3. As the modified model outperforms the Yang-2018 model on all metrics, we use the model proposed by Qi et al. (2019) throughout our experiments as well as the user study.

E Hierarchical Error Propagation

As the S&F model first selects the supporting facts and then predicts the answer based on the selected subset, we evaluate how errors in the fact selection effect the answer prediction and compare it to the Qi-2019 model. For both models, we compare the fraction of predictions with correct (exact match) answers among predictions that (i) contain all ground truth facts or (ii) contain no ground truth facts. We observe that for Qi-2019 the fraction of correct answers drops from 51.5% for predictions with all ground truth facts to 25.8% for predictions with no ground truth facts. For the S&F model, we observe a drop from 50.4% to 20.6% respectively.




Figure 6: Answer changes $c_{\text{rel}}(k)$ and $c_{\text{irr}}(k)$ when removing facts predicted as relevant (top) and irrelevant (bottom).

This confirms our expectations of an increased error propagation in the S&F model. However, only 1.8% of the S&F fact predictions contain no ground truth fact at all, whereas for Qi-2019 this case occurs for 2.1% of the predictions.

F Answer Changes per Removal Step

Figure 6 shows the values of $c_{\text{rel}}(k)$ (top), i.e., the fraction of changed answers when removing k facts predicted as relevant, and $c_{\text{irr}}(k)$ (bottom), i.e., the fraction of changed answers when removing k facts predicted as irrelevant, for $k \in \{0, \dots, 4\}$.

G User Study Design Details

All questions and statements that we ask participants to answer/rate are listed in Table 4. Figure 7 and Figure 8 show screenshots of the study interface for the question rating and the post-questionnaire.

H Detailed User Study Results

Figure 9 shows boxplots per condition for all continuous dependent variables. Figure 10 shows rat-

ing distributions per condition for all ordinal dependent variables. Figures 11 and 12 compare model scores and human measures grouped into F_1 -scores and our proposed FARM(4) and LOCA scores. Rows alternate between F_1 -scores and our scores. Table 5 shows pairwise Pearson correlation coefficients between human and automatized scores.

Question: Are Brainerd Lakes Regional Airport and Sawyer International Airport located in Europe?

System answer: no

System explanation:

- [Brainerd Lakes Regional Airport]:
Brainerd Lakes Regional Airport (IATA: BRD, ICAO: KBRD, FAA LID: BRD) is a public use airport located three nautical miles (6 km) northeast of the central business district of Brainerd, a city in Crow Wing County, Minnesota, United States.
- [Sawyer International Airport]:
Sawyer International Airport (IATA: MQT, ICAO: KSAW, FAA LID: SAW) is a county owned public use airport in Marquette County, Michigan, United States.

Do you think the system's answer is correct?

☐

Yes

☐

No

Please rate how much you disagree / agree to each of the following statements.

I am confident that my choice is correct.

☐

Strongly disagree

☐

Disagree

☐

Neutral

☐

Agree

☐

Strongly agree

The given explanation helps me to decide if the answer is correct.

☐

Strongly disagree

☐

Disagree

☐

Neutral

☐

Agree

☐

Strongly agree

Did you know the answer without the system's answer or explanations?

☐

Yes

☐

No

SUBMIT

Figure 7: Screenshot of the question rating interface.

Context	Question/Statement	Answer Range
Each Question	Do you think the system's answer is correct?	yes/no
	Did you know the answer without the system's answer or explanations?	yes/no
	I am confident that my choice is correct.	5-point Likert
	The given explanation helps me to decide if the answer is correct.	5-point Likert
Post Survey	I trust the question answering system.	5-point Likert
	The explanations contained relevant information.	5-point Likert
	The explanations also contained irrelevant information.	5-point Likert
	I am satisfied with the question answering system and its explanations.	5-point Likert

Table 4: Questions and statements shown to the participants for (a) each question (upper part) and (b) in the post questionnaire (lower part). Statements were presented along with the prompt “Please rate how much you disagree/agree to each of the following statements”.

Questionnaire

The following questions are asked with regard to all model outputs you saw on the previous pages.

Please rate how much you disagree / agree to each of the following statements.

I trust the question answering system.

☐ Strongly disagree

☐ Disagree

☐ Neutral

☐ Agree

☐ Strongly agree

The explanations contained relevant information.

☐ Strongly disagree

☐ Disagree

☐ Neutral

☐ Agree

☐ Strongly agree

The explanations also contained irrelevant information.

☐ Strongly disagree

☐ Disagree

☐ Neutral

☐ Agree

☐ Strongly agree

I am satisfied with the question answering system and its explanation.

☐ Strongly disagree

☐ Disagree

☐ Neutral

☐ Agree

☐ Strongly agree

SUBMIT

Figure 8: Screenshot of the post-questionnaire.

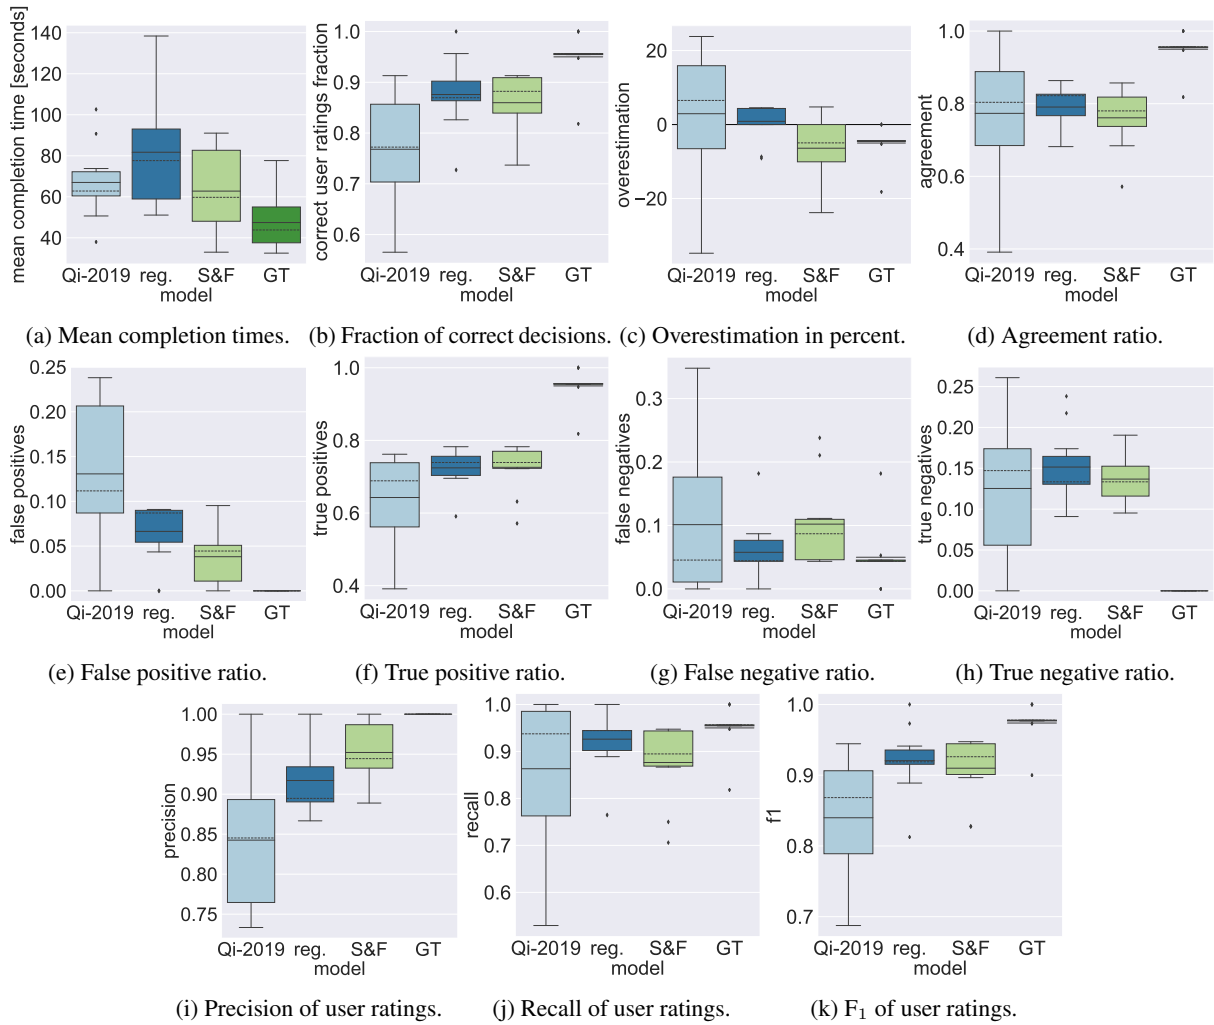


Figure 9: Boxplots for all continuous dependent variables. Boxes mark quartiles, whiskers mark 1.5 inter-quartile ranges, outliers are plotted separately. Vertical solid lines within boxes mark means, vertical dashed lines mark medians.



Figure 10: Distribution of Likert scale ratings. White dots mark number of participants, bar widths correspond to normalized frequency counts. Certainty and helpfulness ratings are aggregated per participant using the participant ratings' mode.

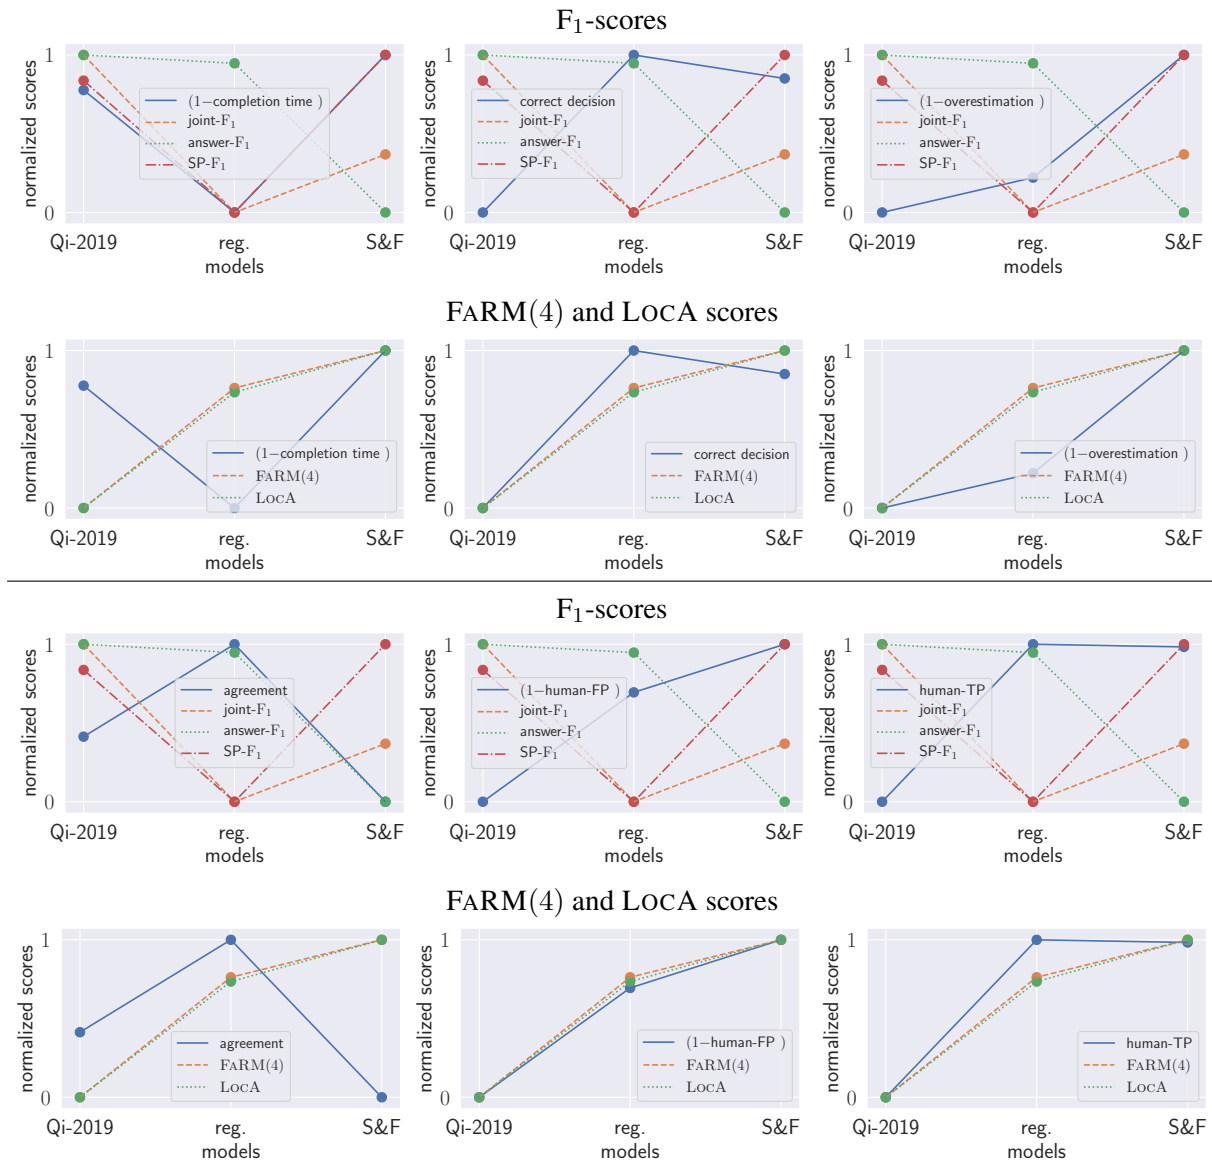


Figure 11: Comparisons between human measures and model scores. All scores are normalized before plotting by subtracting the minimum score and re-scaling the score span to $[0, 1]$. Human measures for which lower values correspond to better performance are plotted as $(1 - \text{score})$ for convenience of the reader. The figure shows scores for completion time, fraction of correct user decisions, overestimation, agreement, false positives and true positives.

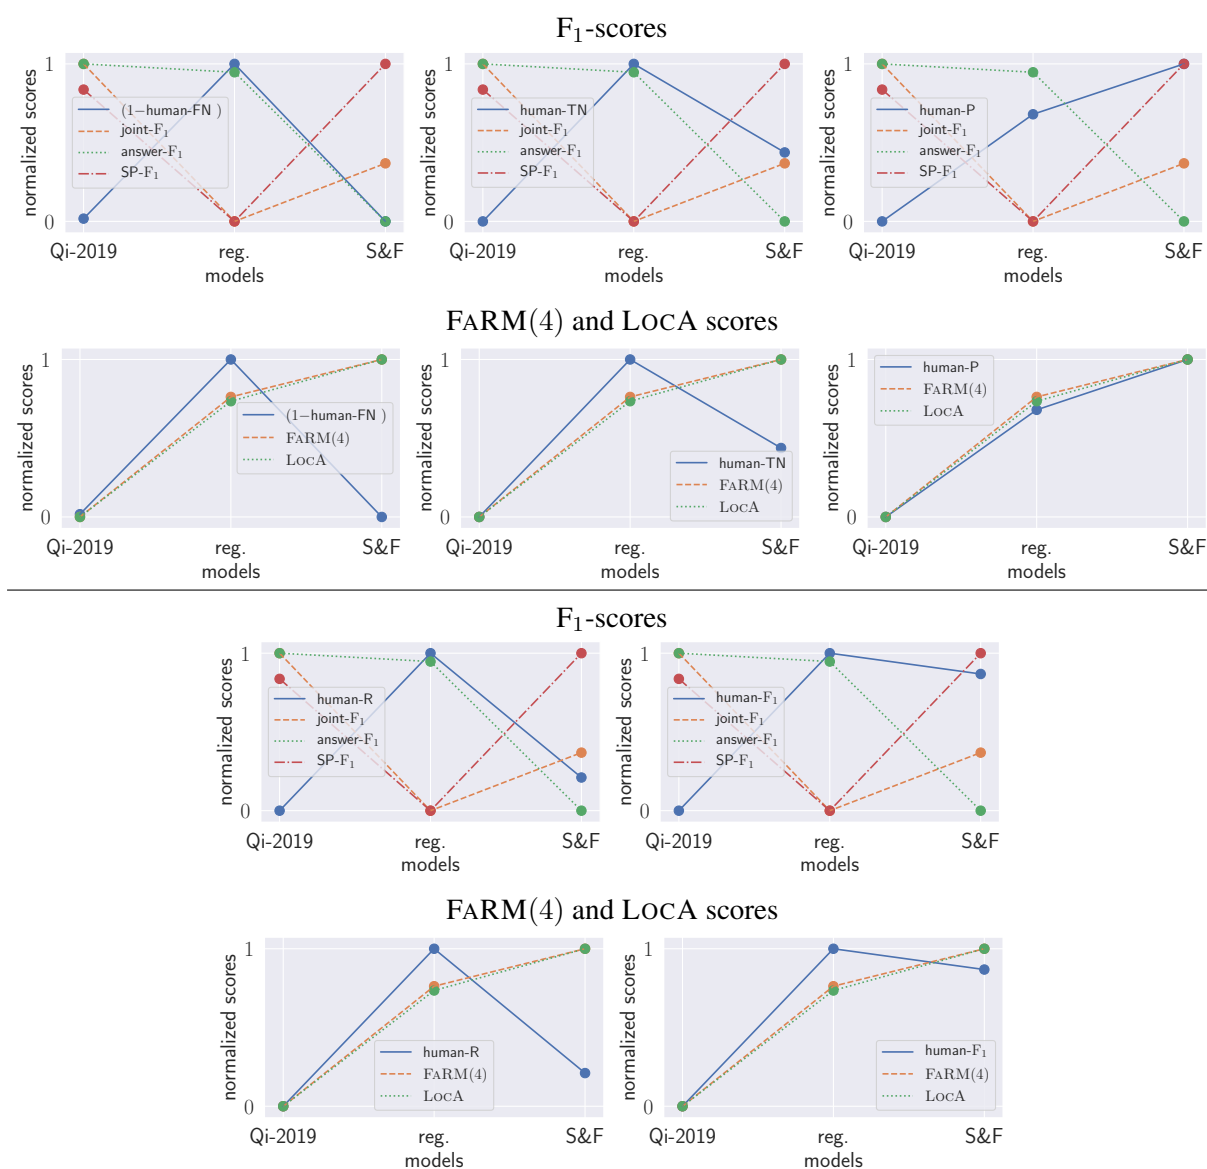


Figure 12: Comparisons between human measures and model scores. All scores are normalized before plotting by subtracting the minimum score and re-scaling the score span to $[0, 1]$. Human measures for which lower values correspond to better performance are plotted as $(1 - \text{score})$ for convenience of the reader. The figure shows scores for false negatives, true negatives, precision, recall and user F_1 .

	Standard Scores												Proposed Scores							
	Answer				Supporting Facts				Joint				Answer Changes		FARM		%in-fact		LOCA	
	EM	F ₁	P	R	EM	F ₁	P	R	EM	F ₁	P	R	rel.	irrel.	FARM(1)	FARM(4)	rel.	irrel.	LOCA	
Human eval.																				
correct decision	-0.33	-0.42	-0.37	-0.47	-0.50	-0.49	-0.57	0.63	-0.57	-0.97	-0.79	0.30	0.80	-0.93	0.99	0.93	1.00	-0.76	0.92	
overestimation	0.97	0.99	0.98	0.99	-0.43	-0.45	-0.36	0.29	-0.36	0.35	-0.07	0.62	0.96	-0.74	0.98	1.00	0.93	-0.94	1.00	
completion time	0.71	0.64	0.67	0.59	-1.00	-1.00	-0.99	0.97	-0.99	-0.63	-0.90	0.99	0.23	-0.95	0.67	0.49	0.79	-0.17	0.46	
human-FP	0.70	0.77	0.73	0.80	0.09	0.07	0.16	-0.24	0.17	0.78	0.45	0.14	0.97	-0.30	0.73	0.87	0.61	-0.98	0.88	
human-TP	-0.45	-0.53	-0.49	-0.58	-0.39	-0.38	-0.46	0.53	-0.47	-0.94	-0.70	0.18	0.87	-0.87	1.00	0.97	0.99	-0.84	0.96	
human-FN	-0.55	-0.47	-0.51	-0.42	0.99	0.99	1.00	-1.00	1.00	0.77	0.97	-0.94	-0.03	0.87	-0.51	-0.30	-0.65	-0.03	-0.28	
human-TN	0.12	0.02	0.07	-0.04	-0.83	-0.82	-0.87	0.91	-0.87	-0.98	-0.98	0.69	0.45	-1.00	0.83	0.68	0.91	-0.40	0.66	
human-P	-0.71	-0.78	-0.75	-0.81	-0.07	-0.05	-0.14	0.22	-0.15	-0.77	-0.43	-0.15	0.98	-0.66	0.95	1.00	0.88	-0.97	1.00	
human-R	0.36	0.27	0.32	0.21	-0.94	-0.94	-0.97	0.98	-0.97	-0.89	-1.00	0.85	0.22	-0.95	0.66	0.48	0.78	-0.16	0.45	
human-F ₁	-0.35	-0.43	-0.39	-0.49	-0.49	-0.47	-0.55	0.62	-0.56	-0.97	-0.78	0.28	0.81	-0.92	0.99	0.94	1.00	-0.77	0.93	

Table 5: Pearson correlations between human and automatized scores.