

Written Justifications are Key to Aggregate Crowdsourced Forecasts

Saketh Kotamraju

Heritage High School

saketh.kotamraju@gmail.com

Eduardo Blanco

Arizona State University

eduardo.blanco@asu.edu

Abstract

This paper demonstrates that aggregating crowdsourced forecasts benefits from modeling the written justifications provided by forecasters. Our experiments show that the majority and weighted vote baselines are competitive, and that the written justifications are beneficial to call a question throughout its life except in the last quarter. We also conduct an error analysis shedding light into the characteristics that make a justification unreliable.

1 Introduction

The wisdom of the crowd refers to the idea that aggregating information collected from many non-experts often yields good answers to questions—as close to the truth or even better than asking an expert. Perhaps the best known example is by Galton (1907), who observed that the median estimate of the weight of an ox (out of 800 country fair attendees) was within 1% of the truth. There is a lot of support for the idea, although it is well known that it is not foolproof (Surowiecki, 2005). MacKay (1841) presents historical examples where crowds behaved irrationally, and more recently, world chess champion Gary Kasparov beat the crowd playing chess (Marko and Haworth, 1999).

In this day and age, the benefits of the crowd are commonplace. Wikipedia is written by volunteers, and community question answering has received the attention of researchers (Adamic et al., 2008; Wang et al., 2013). When aggregating information collected from crowds, it may be important to know whether judgments were collected independently of each other. If they were not, crowd psychology (Reicher, 2001) and the power of persuasion (O’keefe, 2015) can bias individual judgments and eliminate the wisdom of the crowd.

In this paper, we work with forecasts about questions across the political, economic, and social

Question: Will there be a new prime minister of Italy before 1 September 2021?

Start date: 1/28/2021, closing date: 2/13/2021

Forecast 1: 100% yes, 0% no

Justification: *Actually the media talk about potential candidates [link] the Crowd is 98% Yes*

Forecast 2: 99% yes, 1% no

Justification: *With a substantial majority now backing Draghi (who in turns seems to be an obvious EU favourite which brings better prospects for bail out funding) this seems to be a virtual certainty at this stage. [link] Thanks [user] for digging up the parliamentary numbers! [link] [link]*

Figure 1: Question and forecasts submitted by the crowd. Justifications provide information about the credibility of the forecast. The first justification is weak and refers to the current opinion of the crowd; the second justification is strong and provides links to support the claims.

spectrum. Each forecast consists of a prediction estimating the likelihood of some event and a written justification explaining the prediction. As Figure 1 shows, forecasts with the same predictions may come with weaker or stronger justifications that affect the credibility of the predictions. For example, the first justification refers to an external source without justifying why, and it appears to rely on the current opinion of the crowd. On the other hand, the second justification provides specific facts from external resources and previous forecasters.

We move to a discussion of important terminology. We define a *question* as a sentence that elicits its information (e.g., ‘Will legislation raising the US federal minimum wage become law before 20 August 2021?’). Questions have an *opening* and *closing day*, and the days in between are the *life of the question*. *Forecasters* are people who submit a forecast. A *forecast* consists of a prediction and a justification. The *prediction* is a number indicating the chances that something will happen. Following

with the question above, a prediction could be ‘70% chance’ (of the legislation becoming law before 20 August 2021). A *justification* is the text forecasters submit in support of their predictions (see examples in Figure 1 and Section 5). We use the phrase *call a question* to refer to the problem we work with: make a final prediction after aggregating individual forecasts. We call questions each day throughout their life using two strategies: forecasts submitted in the given day (*daily*) and the last forecast submitted by each forecaster (*active*). Note that in this paper we use *prediction* to refer to the submission by a forecaster, not the output of a machine learning model.

Inspired by previous work on identifying and cultivating better forecasters (Mellers et al., 2015), and analyzing written justifications to estimate the quality of a single forecast (Schwartz et al., 2017) or all forecasts by a forecaster (Zong et al., 2020), we experiment with the problem of automatically calling a question through its life based on the available forecasts in each day. The main contributions of this paper are empirical results answering the following research questions:

- When calling a question on a particular day, is it worth taking into account forecasts submitted in previous days? (it is);
- Does calling a question benefit from taking into account the question and the justifications submitted with the forecasts? (it does);
- Is it easier to call a question towards the end of its life? (it is); and
- Is it true that the worse the crowd predictions the more useful the justifications? (it is).

In addition, we also present an analysis of the justifications submitted with correct and wrong forecasts to shed light into which characteristics make a justification more and less credible.

2 Previous Work

The language people use is indicative of several attributes. Previous work includes both predictive models (input: language samples, output: some attribute about the author) and models that yield useful insights (input: language samples and attributes of the authors, output: differentiating language features depending on the attributes). Among many others, previous research has studied gender and age (Li et al., 2016; Nguyen et al., 2014; Peersman et al., 2011), political ideology (Iyyer et al., 2014; Preoțiuc-Pietro et al., 2017), health

outcomes (Schneuwly et al., 2019), and personality traits (Schwartz et al., 2013). In this paper, we do not profile forecasters. Instead, we build models to call questions based on forecasts by the crowd without knowledge of who submitted what.

Previous research has also studied the language people use to communicate depending on the relationship between the parties. For example, the language people use when they are in positions of power (e.g., more seniority) has been studied in social networks (Bramsen et al., 2011), online communities (Danescu-Niculescu-Mizil et al., 2012), and corporate emails (Prabhakaran and Rambow, 2014). Similarly, Rashid and Blanco (2018) study how language provides clues about the interactions and relationships between people. Regarding language form and functions, prior research has analyzed politeness (Danescu-Niculescu-Mizil et al., 2013), empathy (Sharma et al., 2020), advice (Govindarajan et al., 2020), condolences (Zhou and Jurgens, 2020) usefulness (Momeni et al., 2013), and deception (Soldner et al., 2019). More related to the problem we work with, Maki et al. (2017) analyze the influence of Wikipedia editors, and Katerenchuk and Rosenberg (2016) study influence levels in online communities. Persuasion has also been studied from a computational perspective (Wei et al., 2016; Yang et al., 2019), including dialogue systems (Wang et al., 2019). The work presented here complements these works. We are interested in identifying credible justifications in order to aggregate crowdsourced forecasts, and we do so without explicitly targeting any of the above characteristics.

Within computational linguistics, the previous task that is perhaps the closest to our goal is argumentation: a good justification for a forecast is arguably a good supporting argument. Previous work includes identifying argument components such as claims, premises, backings, rebuttals, and refutations (Habernal and Gurevych, 2017), and mining supporting and opposing arguments for a claim (Stab et al., 2018). Notwithstanding these works, we found that crowdsourced justifications rarely fall into these argumentation frameworks despite the former are useful to aggregate forecasts.

Finally, there are a few works on forecasting that use the same or very similar corpora than we do. From a psychology perspective, Mellers et al. (2015) present strategies to improve forecasting accuracy (using top forecasters, i.e., superforecasters)

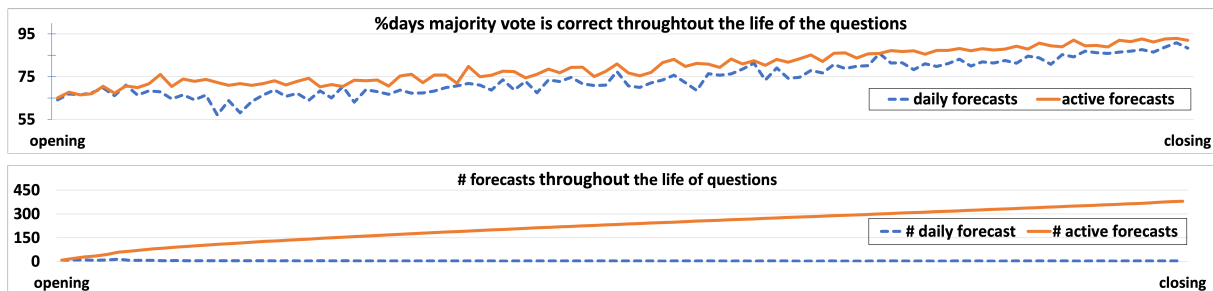


Figure 2: Average number of daily and active forecasts available per question (bottom) and average number of questions the majority forecast gets correct (top) over the life of the question (x-axis). There is a tiny peak of forecasts submitted soon after a question is published and then a roughly uniform amount through the life of the question. The majority of the forecasts, unsurprisingly, is less reliable towards the first half of the life of a question.

and analyze the characteristics of superforecaster performance, which can be used for cultivating better forecasters. Mellers et al. (2014) discuss explanations of what makes forecasters better. These works aim at identifying superforecasters and do not take into account the written justifications. Unlike them, we build models to call questions without using any information about forecasters. Within computational linguistics, Schwartz et al. (2017) assess the language of quality justifications (rating, benefit, and influence). Zong et al. (2020) is perhaps the closest experiment to ours. They build models to predict forecaster skill using the text justifications of forecasts from Good Judgment Open data, and they also use another dataset, Company Earnings Reports, to individually predict which forecasts are more likely to be correct predictions. Unlike us, none of these works aim at calling the question throughout its life.

3 Dataset

We work with data from the Good Judgment Open,¹ a website where questions are posted and people submit forecasts. Questions are about geopolitics and include topics such as domestic and international politics, the economy, and social issues. We collected all binary questions along with all their forecasts including a prediction and a justification. In total, the dataset we work with contains 441 questions and 96,664 forecasts submitted in 32,708 days. This is almost twice the amount of forecasts considered by Zong et al. (2020). Since our goal is to call questions throughout their life, we work with all forecasts with written justifications regardless of length, how many forecasts have been submitted by the same forecaster, etc. Additionally, our

¹<https://www.gjopen.com/>

	Min	Q1	Q2	Q3	Max	Mean
#tokens	8	16	20	28	48	21.94
#entities	0	2	3	5	11	3.47
#verbs	0	2	2	3	6	2.26
#days open	2	24	59	98	475	74.16

Table 1: Analysis of the questions from our dataset. Most questions are relatively long, contain two or more named entities, and are open for over one month.

framework preserves privacy as we do not use any information about the forecaster.

The bottom plot in Figure 2 shows the average number of daily and active forecasts over the life of all questions. There is roughly a uniform number of forecasts submitted each day, thus the amount of active forecasts increases linearly over the life of the question. The majority baseline with both daily and active forecasts submitted in the previous 10 days is quite accurate, especially towards the closing date of questions. The experiments presented in this paper aim at calling questions throughout their life. As we shall see, models to automatically call questions benefit from taking into account justifications during the first three quarters of the life of a question.

Analyzing the Questions Table 1 shows a basic analysis of the questions in our dataset. The majority of questions have over 16 tokens and several entities; the most common are *geopolitical*, *person* and *date* entities. Regarding the life of questions, we observe that half are open for almost two months, and 75% for over three weeks.

Figure 3 shows the LDA topics (Blei et al., 2003) obtained with gensim (Řehůřek and Sojka, 2010). We observe three main topics: elections (voting, winners, candidate, etc.), government actions (ne-




Figure 3: Topics obtained with LDA topic modeling in the 441 questions in our corpus. The topics roughly correspond to (clockwise from top left) (a) elections, (b) government actions, and (c) war and violent events.

	Min	Q1	Q2	Q3	Max
#sentences	1	1	1	3	56
#tokens	1	10	23	47	1295
#entities	0	0	2	4	154
#verbs	0	1	3	6	174
#adverbs	0	0	1	3	63
#adjectives	0	0	2	4	91
#negation	0	0	1	3	69
Sentiment	-2.54	0	0	0.20	6.50
Readability					
Flesch	-49.68	50.33	65.76	80.62	121.22
Dale-Chall	0.05	6.72	7.95	9.20	19.77

Table 2: Analysis of the 96,664 written justifications submitted by forecasters in our dataset. The readability scores indicate that most justifications are easily understood by high school students (11th or 12th grade), although a substantial amount (>25%) require a college education (Flesch under 50 or Dale-Chall over 9.0).

gotiations, announcements, meetings, passing (a law), etc.), and wars and violent crimes (groups killing, civilian (casualties), arms, etc.). While not shown in the LDA topics, the questions cover both domestic and international events in these topics.

Analyzing the Justifications Table 2 presents basic analysis of the 96,664 forecasts justification in our dataset. The median length is short (1 sentence and 23 tokens), and justifications mention named entities less often than questions (Table 1). We check whether justifications have negations using the cues annotated in ConanDoyle-neg (Morante and Daelemans, 2012). Surprisingly, half of the justifications have one negation, and 25% have three or more. This indicates that forecasters sometimes rely on what may *not* happen (or has *not* happened)

to make predictions about the future (questions do not have negations). We also look at the sentiment polarity of justifications using TextBlob (Loria, 2020). The majority of justifications are neutral (polarity close to 0). In terms of readability, we compute the Flesch (Flesch, 1948) and Dale-Chall (Dale and Chall, 1948) scores. Both scores indicate that around a quarter of justifications require a college education to be understood.

In terms of verbs and nouns, we analyze them using the WordNet lexical files (Miller, 1995). The most common verb classes are *change* (26% of justifications, e.g., happen, remain, increase) *social* (24%, e.g., vote, support, help) *cognition* (22%, e.g., think, believe, know) and *motion* (19%, e.g., go, come, leave). The most common noun classes are *act* (71%, e.g., election, support, deal), *communication* (57%, e.g., questions, forecast, news), *cognition* (38%, e.g., point, issue, possibility), and *group* (38%, e.g., government, people, party).

4 Experiments and Results

We experiment with the problem of calling a question throughout its life. The input to the problem is the question itself and forecasts (predictions and justifications), and the output is an answer to the question aggregating all the forecasts. The number of instances is the number of days all questions were open (recall our dataset contains 441 questions and 96,664 forecasts submitted in 32,708 days). We experiment both with simple baselines and a neural network taking into account (a) daily forecasts and (b) active forecasts submitted up to ten days prior. Experimental results showed that considering earlier active forecasts is not beneficial.

We divide the questions into training, validation, and test subsets. Then, we assign to each subset all the forecasts submitted throughout the life of the questions. Note that randomly splitting forecasts would be unsound, as forecasts for the same questions submitted on different days would end in the training, validation, and test subsets.

Baselines We consider two unsupervised baselines. The *majority* vote baseline calls a question based on the majority prediction in the forecasts. The *weighted* vote baseline calls a question after weighting the chances assigned to the predictions in the forecasts. Consider these three forecasts: 99%, 45%, and 45% chance the answer is *yes* (thus 1%, 55%, and 55% chance the answer is *no*). The majority vote baseline would output *no* (2 out of 3 believe

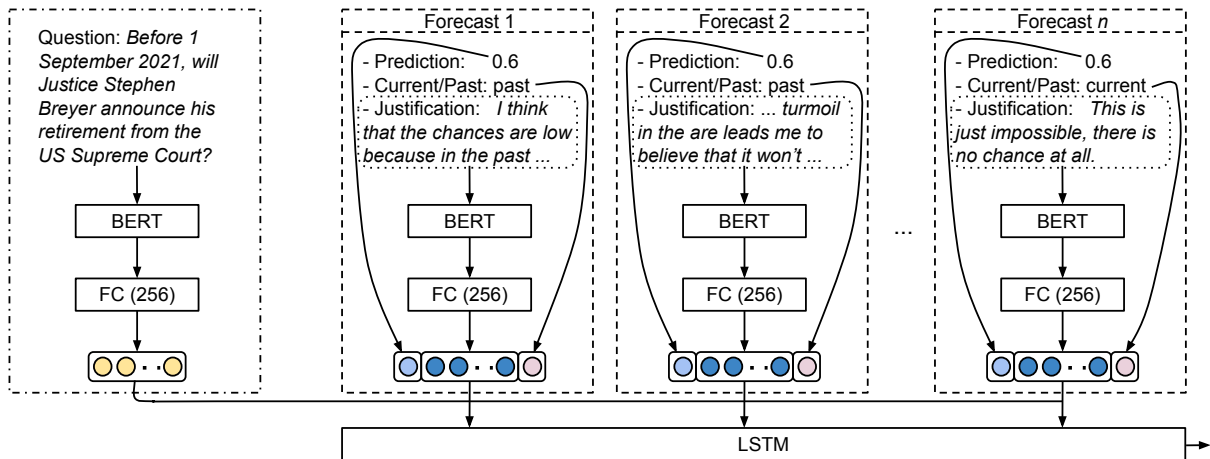


Figure 4: Neural network architecture to call a question on a given day based on crowdsourced forecasts. The network consists of three main components: one for the question, one for each forecast (prediction + flag indicating current day or past + justification), and an LSTM to process the sequence of forecasts. We experiment with two scenarios: feeding the network the forecasts submitted on a given day (daily) or the last forecast by each forecaster within the ten previous days of a given day (active).

no is more likely). On the other hand, the weighted vote baseline would output *yes* (the weighted support for *yes* is larger, 0.99 vs. 0.90).

4.1 Neural Network Architecture

We experiment with the neural network architecture depicted in Figure 4. The network has three main components: a component to obtain a representation of the question, a component to obtain a representation of a forecast, and an LSTM (Hochreiter and Schmidhuber, 1997) to process the sequence of forecasts and call the question.

We obtain the representation of a question using BERT (Devlin et al., 2019) followed by a fully connected layer with 256 neurons, ReLU activation, and 0.5 dropout (Srivastava et al., 2014). We obtain the representation of a forecast concatenating three elements: (a) a binary flag indicating whether the forecast was submitted in the day the question is being called or in the past, (b) the prediction (a number ranging from 0.0 to 1.0), and (c) a representation of the justification. We obtain the representation of the justification using BERT followed by a fully connected layer with 256 neurons, ReLU activation, and 0.5 dropout. The LSTM has a hidden state with dimensionality 256, and takes as its input the sequence of forecasts. During the tuning process, we discovered that it is beneficial to pass the representation of the question with each forecast as opposed to processing forecasts independently of the question. Therefore, we concatenate the representation of the question to each representation

of a forecast prior to feeding the sequence to the LSTM. Finally, the last hidden state of the LSTM is connected to a fully connected layer with 1 neuron and sigmoid activation to call the question.

Architecture Ablation We experiment with the full neural architecture as described above and disabling several components. Specifically, we experiment with representing a forecast taking into account different information:

- the *prediction*;
- the *prediction* and the representation of the *question*;
- the *prediction* and the representation of the *justification*; and
- the *prediction*, the representation of the *question*, and the representation of the *justification*.

Implementation and Training Details In order to implement the models,² we use the Transformers library by HuggingFace (Wolf et al., 2020) and PyTorch (Paszke et al., 2019). We use binary cross-entropy loss, gradient accumulation and mixed precision training (Micikevicius et al., 2018) to alleviate the memory requirements, the Adam optimizer (Kingma and Ba, 2015) with learning rate 0.001, batch size 16, and early stopping with patience set to 3 epochs. We tuned all the hyperparameters comparing held-out results with the validation set, and report results with the test set.

²Code to replicate our experiments available at https://github.com/saketh12/forecasting_emnlp2021

	days with question was open				
	All	Q1	Q2	Q3	Q4
Using daily forecasts only					
Baselines					
Majority vote (predictions)	71.89 [†]	64.59 [†]	66.59 [†]	73.26 [†]	82.22 [†]
Weighted vote (predictions)	73.79 [†]	67.79 [†]	68.71 [†]	74.16 [†]	83.61 [†]
Neural network with components ...					
predictions	77.96 [†]	77.62	77.93	78.23	78.61 [†]
predictions + question	77.61 [†]	75.44 [‡]	76.77 [†]	78.05 [†]	81.56 ^{†‡}
predictions + justifications	80.23 [‡]	77.87	78.65	79.26	84.67 ^{†‡}
predictions + question + justifications	79.96 ^{†‡}	78.65	78.11 [†]	80.29 ^{†‡}	83.28 ^{†‡}
Using active forecasts					
Baselines					
Majority vote (predictions)	77.27 [†]	68.83 [†]	73.92 [†]	77.98 [†]	87.44 [†]
Weighted vote (predictions)	77.97 [†]	72.04 [†]	72.17 [†]	78.53 [†]	88.22 [†]
Neural network with components ...					
predictions	78.81 [†]	77.31	78.04	78.53	81.11 [†]
predictions + question	79.35 [†]	76.05	78.53 [†]	79.56 [†]	82.94 ^{†‡}
predictions + justifications	80.84 [‡]	77.86	79.07	79.74	86.17 ^{†‡}
predictions + question + justifications	81.27 ^{†‡}	78.71 [‡]	79.81 ^{†‡}	81.56 ^{†‡}	84.67 ^{†‡}

Table 3: Results with the test questions (Accuracy, i.e., the average percentage of days a model calls a question correctly). We provide results with *All* days a question was open and four quartiles (Q1: first 25% of days, Q2: 25–50%, Q3: 50–75%, and Q4: last 25% of days). We calculate statistical significance (McNemar’s test (McNemar, 1947) with $p < 0.05$) between (a) each model using daily or active forecasts (all models obtain significantly better results using the active forecasts except the neural network with the *predictions + justifications* component, indicated with [†]) and (b) the neural network trained with the *predictions* component and the networks trained with the additional components (adding the *justification* and both the *question and justification* yields significantly better results using daily or active forecasts, indicated with [‡]).

4.2 Quantitative Results

Table 3 presents the results. The evaluation metrics is accuracy (i.e., average percentage of days a model calls a question correctly throughout the life of the question). We report results for all days (column 2) and the four quartiles (columns 3–6).

Despite their simplicity, the baselines obtain good results (71.89 and 73.79 using daily and active forecasts), showing that aggregating the predictions submitted by forecasters without regard to the question or justifications is a competitive approach. As we shall see, however, the full neural network obtains statistically significant better results (79.96 and 81.27 using daily and active forecasts).

Using Daily or Active Forecasts Taking into account active forecasts instead of only those submitted on the day the model is calling the question (daily forecasts) is beneficial across both baselines and all neural networks except the one using only *predictions + justification*. The differences in accuracy are larger with the baselines (daily: 71.89 vs. 77.27; active: 73.79 vs. 77.97) than with the

neural networks. We note, however, that the differences are statistically significant evaluating with all days and all quartiles except Q1 (indicated with [†] in Table 3, McNemar’s test (McNemar, 1947) with $p < 0.05$). **We conclude that using active forecasts is beneficial** and focus the remaining of the discussion on these results.

Encoding Questions and Justifications The neural network that uses only the *prediction* to represent a forecast outperforms both baselines (78.81 vs. 77.27 and 77.97). More interestingly, incorporating into the representation of the forecast the question, the justification, or both brings improvements (79.35, 80.84, and 81.27). All but the results with *predictions + justifications* are statistically significant with respect to using only *predictions*. **We conclude that calling a question benefits from incorporating into the model the question and the justifications submitted by forecasters.**

Calling Questions Throughout their Life We now move beyond accuracies calculated using all days throughout the life of a question and exam-

	question difficulty (according to best baseline)				
	All	Q1	Q2	Q3	Q4
Using active forecasts					
Weighted vote baseline (predictions)	77.97	99.40	99.55	86.01	29.30
Neural network with components ...					
predictions + question	79.35 [†]	94.58 ^{†‡}	88.01 ^{†‡}	78.04 ^{†‡}	58.73 ^{†‡}
predictions + justifications	80.84 ^{†‡}	95.71 ^{†‡}	93.18 ^{†‡}	79.99 ^{†‡}	57.05 ^{†‡}
predictions + question + justifications	81.27 ^{†‡}	94.17 ^{†‡}	90.11 ^{†‡}	78.67 ^{†‡}	64.41 ^{†‡}

Table 4: Results with the test questions (Accuracy, i.e., average percentage of days a question is called correctly). We provide results with *All* questions and depending on the question difficulty as measured by the results obtained with the best baseline (Q1: easiest 25%; Q2: 25–50%, Q3: 50–75%, and Q4: hardest 25%). We calculate statistical significance (McNemar’s test (McNemar, 1947) with $p < 0.05$) between (a) the weighted vote baseline and each neural network (indicated with [†]), and (b) the neural network trained with the *predictions* component (not shown) and the networks trained with the additional components (indicated with [‡]).

ine detailed results per quartile. More specifically, we divide the days into four quartiles. The last four columns in Table 3 show that while using active forecasts is beneficial across all four quartiles (with both baselines and all networks), the neural networks—perhaps surprisingly—outperform the baselines only in the first three quartiles. In fact, the neural networks obtain statistically significant worse results than any of the baselines in the last quartile (84.67 vs. 87.44 and 88.22; -3.2% and -4.0%). **We conclude that modeling questions and justifications is overall useful, although it is detrimental towards the end of the life of a question.** The justification for this empirical fact is that the crowd gets wiser towards the end of the life of a question—as more evidence to make the correct prediction presumably becomes available, and more forecasters submit forecasts. Our model does not take into account which day is calling a question in (within the life of the question). We reserve to future work incorporating temporal information to better aggregate forecasts.

Calling Questions Based on their Difficulty We finish the quantitative experiments with results depending on the difficulty of the questions. To this end, we sort questions by their difficulty based on how many days the majority or weighted vote baselines (whichever makes the least mistakes) calls the questions wrong. These experiments shed light into how many questions benefit from the neural networks that take into account the question and justifications. We note, however, that it is impossible to calculate question difficulty during the life of the question, so these experiments are not realistic before a question closes (and the correct answer is

known). After all, forecasts are about predicting the future, and it is only challenging to do so while the correct answer is unknown.

Table 4 shows the results with selected models depending on question difficulty. We observe that the weighted vote baseline calls 75% questions more reliably than the neural network. Indeed, the baseline obtains 99.40, 99.55, 86.01 and 29.30 accuracy in each quartile of difficulty, while the best network obtains 95.71 (-3.7%), 93.18 (-6.4%), 79.99 (-7.0%), and 64.41 (+119.8%). In other words, the majority of questions (75% easiest questions) obtain worse results with the best neural network (-3.7–7.0%), but a substantial amount (25% hardest questions) are called correctly more than twice as often (+119.8%). The benefits with the hardest questions compensate the drawbacks with the easiest questions. As stated earlier, overall the full neural network obtains significantly better results than the baselines (81.27 vs. 77.27 and 77.97). **We conclude that learning how to aggregate crowdsourced forecasts, and specifically taking into account the question and justifications, is the most beneficial with the hardest questions.**

5 Qualitative Analysis

In this section, we present insights into (a) what makes questions harder to forecast and (b) characteristics of justification submitted with wrong and correct predictions (Table 5).

Questions We looked at three characteristics of the 88 questions in the test set depending on whether the best model (bottom row in Table 3) calls the question at least one day wrong (it does so with 36 out of 88 questions). Surprisingly, we

	Questions . . .	
	called wrong ≥ 1 day	always called correct
# days open	69.4	81.7
# forecasts available	31.0	26.7
% incorrect forecasts	49.7	16.6
	Justifications submitted with . . .	
	wrong predictions	correct predictions
% short (< 20 tokens)	78.0	65.0
% with references to previous forecasts	31.5	16.0
% without a logical argument	62.5	47.5
% with generic arguments	16.0	14.5
% with poor grammar or spelling, non-English	24.5	14.5

Table 5: Characterizations of questions and justifications based on the predictions obtained with the best model (NN with *predictions + question + justification* trained with active forecasts, Table 3). The top block characterizes all questions in the test set (88 questions) depending on whether the model calls the question wrong in one day. The bottom block characterizes 400 random justifications from days that the model calls a question wrong (200 written justifications submitted with correct and wrong forecasts each).

found that questions that are called correct in all days have a longer life (81.7 vs. 69.4 days) and less active forecasts per day (26.7 vs. 31.0). As one would expect, our best model makes mistakes with the same questions that forecasters struggle with.

Justifications. We manually analyzed 200 justifications submitted with wrong and correct predictions (400 in total). Specifically, we looked at predictions submitted on days that our best model makes a mistake calling the corresponding question. Here are the observations we identified:

- We found that 78% of wrong predictions were submitted with *short justifications* (less than 20 tokens), while 65% of correct predictions were. This observation corroborates that longer user-generated text has higher quality (Beygelzimer et al., 2015). Example: *Software isn't good enough yet*, submitted to question *Will Google's AlphaGo beat world champion Lee Sedol in the five-game Go match planned for March 2016?*
- While relatively few forecasts *refer to previous forecasts* (by the same or other forecasters, or the current forecast by the crowd), we observe that justifications for wrong predictions do almost twice as often (31.5% vs. 16.0%). Example: *Returning to initial forecast*.
- *Lack of logical arguments* is common in the justifications we work with. This is true regardless of whether the predictions they were submitted with are wrong or correct. We

found, however, that not having a logical argument is more common with wrong predictions (62.5% vs. 47.5%).

Example: *I guess Greek head of state does not count, but we are getting close*, submitted to question *Will Iran host a head of state or government from one of the G7 countries on an official visit before 1 July 2016?*

- Surprisingly, justifications with *generic arguments* are not a clear indicator of wrong or correct predictions (16.0% vs. 14.5%).

Example: *It seems to be pretty much decided, unless something completely out of the blue happens*.

- *Poor grammar and spelling or non-English* are rare, but much more common in justification of wrong predictions (24.5% vs. 14.5%). Example: *For reference y'all and Wenn Trump den Kurs beibehAlt*.

6 Conclusions

Forecasting is the process of predicting future events. Government and industry alike are interested in forecasting because it affords them the capability to anticipate and address potential challenges to come. In this paper, we work with questions across the political, economic, and social spectrum published in the Good Judgment Open website, and forecasts submitted by the crowd without special training. Each forecast consists of a prediction and a justification in natural language.

We have shown that aggregating the weighted predictions of forecasters is a robust baseline to call a question through its life. Models that take into account both the question and justifications, however, obtain significantly better results when calling a question in the first three quartiles of its life. Crucially, our models do not profile forecasters or use any information about who submitted which forecast. The work presented here opens the door to assessing the credibility of anonymous forecasts in order to come up with aggregation strategies that are robust without tracking forecasters.

Acknowledgments

We gratefully acknowledge the support of NVIDIA Corporation with the donation of the Titan Xp GPU used for this research. The results presented in this paper were also obtained using the Chameleon testbed supported by the National Science Foundation (Keahey et al., 2020).

References

- Lada A Adamic, Jun Zhang, Eytan Bakshy, and Mark S Ackerman. 2008. Knowledge sharing and yahoo answers: everyone knows something. In *Proceedings of the 17th international conference on World Wide Web*, pages 665–674.
- Alina Beygelzimer, Ruggiero Cavallo, and Joel Tetreault. 2015. On yahoo answers, long answers are best. In *CrowdML: The ICML 15 Workshop on Crowdsourcing and Machine Learning*.
- David M Blei, Andrew Y Ng, and Michael I Jordan. 2003. Latent dirichlet allocation. *Journal of Machine Learning Research*, 3(Jan):993–1022.
- Philip Bramsen, Martha Escobar-Molano, Ami Patel, and Rafael Alonso. 2011. [Extracting social power relationships from natural language](#). In *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies*, pages 773–782, Portland, Oregon, USA. Association for Computational Linguistics.
- Edgar Dale and Jeanne S Chall. 1948. A formula for predicting readability: Instructions. *Educational research bulletin*, pages 37–54.
- Cristian Danescu-Niculescu-Mizil, Lillian Lee, Bo Pang, and Jon Kleinberg. 2012. [Echoes of power: Language effects and power differences in social interaction](#). In *Proceedings of the 21st International Conference on World Wide Web, WWW '12*, page 699–708, New York, NY, USA. Association for Computing Machinery.
- Cristian Danescu-Niculescu-Mizil, Moritz Sudhof, Dan Jurafsky, Jure Leskovec, and Christopher Potts. 2013. [A computational approach to politeness with application to social factors](#). In *Proceedings of the 51st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 250–259, Sofia, Bulgaria. Association for Computational Linguistics.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [BERT: Pre-training of deep bidirectional transformers for language understanding](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.
- Rudolph Flesch. 1948. A new readability yardstick. *Journal of applied psychology*, 32(3):221.
- F. Galton. 1907. [Vox populi](#). *Nature*, 75(1949):7.
- Venkata Subrahmanyam Govindarajan, Benjamin Chen, Rebecca Warholc, Katrin Erk, and Junyi Jessy Li. 2020. [Help! need advice on identifying advice](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 5295–5306, Online. Association for Computational Linguistics.
- Ivan Habernal and Iryna Gurevych. 2017. [Argumentation mining in user-generated web discourse](#). *Computational Linguistics*, 43(1):125–179.
- Sepp Hochreiter and Jürgen Schmidhuber. 1997. Long short-term memory. *Neural computation*, 9(8):1735–1780.
- Mohit Iyyer, Peter Enns, Jordan Boyd-Graber, and Philip Resnik. 2014. [Political ideology detection using recursive neural networks](#). In *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1113–1122, Baltimore, Maryland. Association for Computational Linguistics.
- Denys Katerenchuk and Andrew Rosenberg. 2016. Hierarchy prediction in online communities. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 30.
- Kate Keahey, Jason Anderson, Zhuo Zhen, Pierre Riteau, Paul Ruth, Dan Stanzione, Mert Cevik, Jacob Colleran, Haryadi S. Gunawi, Cody Hammock, Joe Mambretti, Alexander Barnes, François Halbach, Alex Rocha, and Joe Stubbs. 2020. Lessons learned from the chameleon testbed. In *Proceedings of the 2020 USENIX Annual Technical Conference (USENIX ATC '20)*. USENIX Association.
- Diederik P. Kingma and Jimmy Ba. 2015. [Adam: A method for stochastic optimization](#). In *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*.

- Shoushan Li, Bin Dai, Zhengxian Gong, and Guodong Zhou. 2016. [Semi-supervised gender classification with joint textual and social modeling](#). In *Proceedings of COLING 2016, the 26th International Conference on Computational Linguistics: Technical Papers*, pages 2092–2100, Osaka, Japan. The COLING 2016 Organizing Committee.
- Steven Loria. 2020. TextBlob: Simplified Text Processing. <https://github.com/sloria/TextBlob>.
- Charles MacKay. 1841. *Memoirs of Extraordinary Popular Delusions and the Madness Of Crowds*. London: Richard Bentley.
- Keith Maki, Michael Yoder, Yohan Jo, and Carolyn Rosé. 2017. [Roles and success in Wikipedia talk pages: Identifying latent patterns of behavior](#). In *Proceedings of the Eighth International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 1026–1035, Taipei, Taiwan. Asian Federation of Natural Language Processing.
- P Marko and Guy McCrossan Haworth. 1999. The kasparov-world match. *ICGA Journal*, 22(4):236–238.
- Quinn McNemar. 1947. Note on the sampling error of the difference between correlated proportions or percentages. *Psychometrika*, 12(2):153–157.
- Barbara Mellers, Eric Stone, Terry Murray, Angela Minster, Nick Rohrbaugh, Michael Bishop, Eva Chen, Joshua Baker, Yuan Hou, Michael Horowitz, et al. 2015. Identifying and cultivating superforecasters as a method of improving probabilistic predictions. *Perspectives on Psychological Science*, 10(3):267–281.
- Barbara Mellers, Lyle Ungar, Jonathan Baron, Jaime Ramos, Burcu Gurcay, Katrina Fincher, Sydney E Scott, Don Moore, Pavel Atanasov, Samuel A Swift, et al. 2014. Psychological Strategies for Winning a Geopolitical Forecasting Tournament. *Psychological Science*, 25(5):1106–1115.
- Paulius Micikevicius, Sharan Narang, Jonah Alben, Gregory Diamos, Erich Elsen, David Garcia, Boris Ginsburg, Michael Houston, Oleksii Kuchaiev, Ganesh Venkatesh, and Hao Wu. 2018. [Mixed precision training](#). In *International Conference on Learning Representations*.
- George A Miller. 1995. Wordnet: a lexical database for english. *Communications of the ACM*, 38(11):39–41.
- Elaheh Momeni, Claire Cardie, and Myle Ott. 2013. Properties, prediction, and prevalence of useful user-generated comments for descriptive annotation of social media objects. In *Proceedings of the International AAAI Conference on Web and Social Media*, volume 7.
- Roser Morante and Walter Daelemans. 2012. [ConanDoyle-neg: Annotation of negation cues and their scope in conan doyle stories](#). In *Proceedings of the Eighth International Conference on Language Resources and Evaluation (LREC'12)*, pages 1563–1568, Istanbul, Turkey. European Language Resources Association (ELRA).
- Dong Nguyen, Dolf Trieschnigg, A. Seza Dođruöz, Rilana Gravel, Mariët Theune, Theo Meder, and Franciska de Jong. 2014. [Why gender and age prediction from tweets is hard: Lessons from a crowdsourcing experiment](#). In *Proceedings of COLING 2014, the 25th International Conference on Computational Linguistics: Technical Papers*, pages 1950–1961, Dublin, Ireland. Dublin City University and Association for Computational Linguistics.
- Daniel J O’keefe. 2015. *Persuasion: Theory and research*. Sage Publications.
- Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Kopf, Edward Yang, Zachary DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. 2019. [Pytorch: An imperative style, high-performance deep learning library](#). In *Advances in Neural Information Processing Systems 32*, pages 8024–8035. Curran Associates, Inc.
- Claudia Peersman, Walter Daelemans, and Leona Van Vaerenbergh. 2011. [Predicting age and gender in online social networks](#). In *Proceedings of the 3rd International Workshop on Search and Mining User-Generated Contents, SMUC ’11*, page 37–44, New York, NY, USA. Association for Computing Machinery.
- Vinodkumar Prabhakaran and Owen Rambow. 2014. [Predicting power relations between participants in written dialog from a single thread](#). In *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 339–344, Baltimore, Maryland. Association for Computational Linguistics.
- Daniel Preotjiuc-Pietro, Ye Liu, Daniel Hopkins, and Lyle Ungar. 2017. [Beyond binary labels: Political ideology prediction of Twitter users](#). In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 729–740, Vancouver, Canada. Association for Computational Linguistics.
- Farzana Rashid and Eduardo Blanco. 2018. [Characterizing interactions and relationships between people](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 4395–4404, Brussels, Belgium. Association for Computational Linguistics.
- Radim Řehůřek and Petr Sojka. 2010. Software Framework for Topic Modelling with Large Corpora. In *Proceedings of the LREC 2010 Workshop on New Challenges for NLP Frameworks*, pages 45–50, Valletta, Malta. ELRA.

- Stephen Reicher. 2001. *The Psychology of Crowd Dynamics*, chapter 8. John Wiley & Sons, Ltd.
- Arno Schnewly, Ralf Grubenmann, Séverine Rion Lojean, Mark Cieliebak, and Martin Jaggi. 2019. [Correlating Twitter language with community-level health outcomes](#). In *Proceedings of the Fourth Social Media Mining for Health Applications (#SMM4H) Workshop & Shared Task*, pages 71–78, Florence, Italy. Association for Computational Linguistics.
- H. Andrew Schwartz, Johannes C. Eichstaedt, Margaret L. Kern, Lukasz Dziurzynski, Stephanie M. Ramones, Megha Agrawal, Achal Shah, Michal Kosinski, David Stillwell, Martin E. P. Seligman, and Lyle H. Ungar. 2013. [Personality, gender, and age in the language of social media: The open-vocabulary approach](#). *PLOS ONE*, 8(9):1–16.
- H. Andrew Schwartz, Masoud Rouhizadeh, Michael Bishop, Philip Tetlock, Barbara Mellers, and Lyle Ungar. 2017. [Assessing objective recommendation quality through political forecasting](#). In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 2348–2357, Copenhagen, Denmark. Association for Computational Linguistics.
- Ashish Sharma, Adam Miner, David Atkins, and Tim Althoff. 2020. [A computational approach to understanding empathy expressed in text-based mental health support](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 5263–5276, Online. Association for Computational Linguistics.
- Felix Soldner, Verónica Pérez-Rosas, and Rada Mihalcea. 2019. [Box of lies: Multimodal deception detection in dialogues](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 1768–1777, Minneapolis, Minnesota. Association for Computational Linguistics.
- Nitish Srivastava, Geoffrey Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov. 2014. [Dropout: A Simple Way to Prevent Neural Networks from Overfitting](#). *Journal of Machine Learning Research*, 15(56):1929–1958.
- Christian Stab, Tristan Miller, Benjamin Schiller, Pranav Rai, and Iryna Gurevych. 2018. [Cross-topic argument mining from heterogeneous sources](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 3664–3674, Brussels, Belgium. Association for Computational Linguistics.
- James Surowiecki. 2005. *The Wisdom of Crowds*. Anchor.
- Gang Wang, Konark Gill, Manish Mohanlal, Haitao Zheng, and Ben Y Zhao. 2013. [Wisdom in the social crowd: an analysis of quora](#). In *Proceedings of the 22nd international conference on World Wide Web*, pages 1341–1352.
- Xuewei Wang, Weiyan Shi, Richard Kim, Yoojung Oh, Sijia Yang, Jingwen Zhang, and Zhou Yu. 2019. [Persuasion for good: Towards a personalized persuasive dialogue system for social good](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 5635–5649, Florence, Italy. Association for Computational Linguistics.
- Zhongyu Wei, Yang Liu, and Yi Li. 2016. [Is this post persuasive? ranking argumentative comments in online forum](#). In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 195–200, Berlin, Germany. Association for Computational Linguistics.
- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Rémi Louf, Morgan Funtowicz, Joe Davison, Sam Shleifer, Patrick von Platen, Clara Ma, Yacine Jernite, Julien Plu, Canwen Xu, Teven Le Scao, Sylvain Gugger, Mariama Drame, Quentin Lhoest, and Alexander M. Rush. 2020. [Transformers: State-of-the-art natural language processing](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 38–45, Online. Association for Computational Linguistics.
- Diyi Yang, Jiaao Chen, Zichao Yang, Dan Jurafsky, and Eduard Hovy. 2019. [Let’s make your request more persuasive: Modeling persuasive strategies via semi-supervised neural nets on crowdfunding platforms](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 3620–3630, Minneapolis, Minnesota. Association for Computational Linguistics.
- Naitian Zhou and David Jurgens. 2020. [Condolence and empathy in online communities](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 609–626, Online. Association for Computational Linguistics.
- Shi Zong, Alan Ritter, and Eduard Hovy. 2020. [Measuring forecasting skill from text](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 5317–5331, Online. Association for Computational Linguistics.