

# CITE: A Corpus of Image–Text Discourse Relations

**Malihe Alikhani**

Computer Science  
Rutgers University

malihe.alikhani@rutgers.edu

**Gerard de Melo**

Computer Science  
Rutgers University

gerard.demelo@rutgers.edu

**Sreyasi Nag Chowdhury**

Max Planck Institute  
for Informatics

sreyasi@mpi-inf.mpg.de

**Matthew Stone**

Computer Science  
Rutgers University

matthew.stone@rutgers.edu

## Abstract

This paper presents a novel crowd-sourced resource for multimodal discourse: our resource characterizes inferences in image–text contexts in the domain of cooking recipes in the form of coherence relations. Like previous corpora annotating discourse structure between text arguments, such as the Penn Discourse Treebank, our new corpus aids in establishing a better understanding of natural communication and common-sense reasoning, while our findings have implications for a wide range of applications, such as understanding and generation of multimodal documents.

## 1 Introduction

“*Sometimes a picture is worth the proverbial thousand words; sometimes a few well-chosen words are far more effective than a picture*” – [Feiner and McKeown \(1991\)](#). Modeling how visual and linguistic information can jointly contribute to coherent and effective communication is a longstanding open problem with implications across cognitive science. As [Feiner and McKeown \(1991\)](#) already observe, it is particularly important for automating the understanding and generation of text–image presentations. Theoretical models have suggested that images and text fit together into integrated presentations via *coherence relations* that are analogous to those that connect text spans in discourse; see [Alikhani and Stone \(2018a\)](#) and Section 2. This paper follows up this theoretical perspective through systematic corpus investigation.

We are inspired by research on text discourse, which has led to large-scale corpora with information about discourse structure and discourse semantics. The Penn Discourse Treebank (PDTB) is one of the most well-known examples ([Mitsakaki et al., 2004](#); [Prasad et al., 2008](#)). However, although multimodal corpora increasingly include

discourse relations between linguistic and non-linguistic contributions, particularly for utterances and other events in dialogue ([Cuayáhuil et al., 2015](#); [Hunter et al., 2015](#)), to date there has existed no dataset describing the coherence of text–image presentations. In this paper, we describe the construction of an annotated corpus that fills this gap, and report initial analyses of the communicative inferences that connect text and accompanying images in this corpus.

As we describe in Section 2, our approach asks annotators to identify the presence of specific inferences linking text and images, rather than to use a taxonomy of coherence relations. This enables us to deal with the distinctive discourse contributions of photographic imagery. We describe our data collection process in Section 3, showing that our annotation scheme allows us to get reliable labels by crowdsourcing. We present analyses in Section 4 that show that our annotation highlights a range of cases where text and images work together in distinctive and theoretically challenging ways, and discuss the implications of our work for the understanding and generation of multimodal documents. We conclude in Section 5 with a number of problems for future research.

## 2 Discourse Coherence and Text–Image Presentations

We begin with an example to motivate our approach and clarify its relationship to previous work. Figure 1 shows two steps in an online recipe for a ravioli casserole from the RecipeQA data set ([Yagcioglu et al., 2018](#)). The image of Figure 1a shows a moment towards the end of carrying out the “covering” action of the accompanying text; that of Figure 1b shows one instance of the result of the “spooning” actions of the text. Cognitive scientists have argued that such images are



(a) TEXT: Cover with a single layer of ravioli. (b) TEXT: Let cool 5 minutes before spooning onto individual plates.

Figure 1: Two steps in a recipe from [Yagcioglu et al. \(2018\)](#) illustrating diverse inferential relationships between text and accompanying imagery in instructions. The recipe is from Autodesk Inc. [www.instructables.com](http://www.instructables.com) and is contributed by [www.RealSimple.com](http://www.RealSimple.com).

much like text contributions in the way their interpretation connects to the broader discourse. In particular, inferences analogous to those used to interpret text seem to be necessary with such images to recognize their spatio-temporal perspective ([Cumming et al., 2017](#)), the objects they depict ([Abusch, 2013](#)), and their place in the arc of narrative progression ([McCloud, 1993](#); [Cohn, 2013](#)). In fact, such inferences seem to be a general feature of multimodal communication, applying also in the coherent relationships of utterance to co-speech gesture ([Lascarides and Stone, 2009](#)) or the coherent relationships of elements in diagrams ([Alikhani and Stone, 2018b](#); [Hiippala and Orekhova, 2018](#)).

In empirical analyses of text corpora, researchers in projects such as the Penn Discourse Treebank ([Miltsakaki et al., 2004](#); [Prasad et al., 2008](#)) have been successful at documenting such effects by annotating discourse structure and discourse semantics via coherence relations. We would like to apply a similar strategy to text-image documents like that shown in Figure 1. However, existing discourse annotation guidelines depend on the distinctive ways that coherence is signaled in text. In text, we find syntactic devices such as structural parallelism, semantic devices such as negation, and pragmatic elements such as discourse connectives, all of which can help annotators to recognize coherence relations in text. Images lack such features. At the same time, characterizing the communicative role of imagery, par-

ticularly photographic imagery, involves a special problem: distinguishing the content that the author specifically aimed to depict from merely incidental details that happen to appear in the scene ([Stone and Stojnic, 2015](#)).

Thus, rather than start from a taxonomy of discourse relations like that used in PDTB, we characterize the different kinds of inferential relationships involved in interpreting imagery separately.

- To characterize temporal relationships between imagery and text, we ask if the image gives information about the preparation, execution or results of the accompanying step.
- To characterize the logical relationship of imagery to text, we ask if the image shows one of several actions described in the text, and if it depicts an action that needs to be repeated.
- To characterize the significance of incidental detail, we ask a range of further questions (some relevant specifically to our domain of instructions), asking about what the image depicts from the text, what it leaves out from the text, and what it adds to the text.

Our approach is designed to elicit judgments that crowd workers can provide quickly and reliably.

This approach allows us to highlight a number of common patterns that we can think of as prototypical coherence relations between images and text. Figure 1a, for example, instantiates a natural **Depiction** relation: the image shows the action described in the text in progress; the mechanics of the action are fully visible in the image, but the significant details in the imagery are all reported in the text as well. Our approach also lets us recognize more sophisticated inferential relationships, like the fact that Figure 1b shows an **Example:Result** of the accompanying instruction. Many of the relationships that emerge from our annotation effort involve newly-identified features of text-image presentations that deserve further investigation: particularly, the use of loosely-related imagery to provide background and motivation for a multimodal presentation as a whole, and depictions of action that seem simultaneously to give key information about the context, manner and result of an action.

### 3 Annotation Effort<sup>1</sup>

Work on text has found that text genre heavily influences both the kinds of discourse relations one finds in a corpus and the way those relations are signalled (Webber, 2009). Since our focus is on developing methodology for consistent annotation, we therefore choose to work within a single genre. We selected instructional text because of its concrete, practical subject matter and because of its step-by-step organization, which makes it possible to automatically group together short segments of related text and imagery.

**Text–Image Pairs.** We base our data collection on an existing instructional dataset, RecipeQA (Yagcioglu et al., 2018). This is the only publicly available large-scale dataset of multimodal instructions. It consists of multimodal recipes—textual instructions accompanied by one or more images.

We excluded documents that either have multiple steps without images or that have multiple images per set. This was so that we could more easily study the direct relationship between an image and the associated text. There are 1,690 documents with this characteristic in the RecipeQA train set. To avoid overwhelming crowd workers, we further filtered those to retain only recipes with 70 or fewer words per step, for a final count of 516 documents (2,047 image–text pairs).

**Protocol.** We recruit participants through Amazon Mechanical Turk. All subjects were US citizens, agreed to a consent form approved by Rutgers’s institutional review board, and were compensated at an estimated rate of USD 15 an hour.

**Experiment Interface.** Given an image and the corresponding textual instruction from the dataset, participants were requested to answer the following 10 questions.

For Question 1, participants were asked to highlight the relevant part of the text. For the others, we solicited True/False responses.

1. Highlight the part of the text that is most related to the image.
2. The image gives visual information about the step described in the text.

3. You need to see the image in order to be able to carry out the step properly.
4. The text provides specific quantities (amounts, measurements, etc.) that you would not know just by looking at the picture.
5. The image shows a tool used in the step but *not* mentioned in the text.
6. The image shows how to prepare before carrying out the step.
7. The image shows the results of the action that is described in the text.
8. The image depicts an action in progress that is described in the text.
9. The text describes several different actions but the image only depicts one.
10. One would have to repeat the action shown in the image many times in order to complete this step.

The interface is designed such that if the answer to Question 8 is TRUE, the subject will be prompted with Question 9 and 10. Otherwise, Question 8 is the last question in the list.

**Agreement.** To assess the inter-rater agreement, we determine Cohen’s  $\kappa$  and Fleiss’s  $\kappa$  values. For Cohen’s  $\kappa$ , we randomly selected 150 image–text pairs and assigned each to two participants, obtaining a Cohen’s  $\kappa$  of 0.844, which indicates almost perfect agreement. For Fleiss’s  $\kappa$  (Fleiss and Cohen, 1973; Cocos et al., 2015; Banerjee et al., 1999), we randomly selected 50 text–image pairs, assigned them to five subjects, and computed the average  $\kappa$ . We obtain a score of 0.736, which indicates substantial agreement (Viera et al., 2005).

### 4 Analysis

**Overall Statistics.** Table 1 shows the rates of true answers for questions Q2–Q10. Subjects reported that in 17% of cases the images did not give any information about the step described in the accompanying text. Such images deserve further investigation to characterize their interpretive relationship to the document as a whole. Our anecdotal experience is that such images sometimes provide context for the recipe, which may suggest that imagery, like real-world events (Hunter et al., 2015), creates more flexible discourse structures than linguistic segments on their own.

Subjects reported that the image was required in order to carry out the instruction only for 6%

<sup>1</sup>The dataset and the code for the machine learning experiments are available at <https://github.com/malihealikhani/CITE>

	Q2	Q3	Q4	Q5	Q6	Q7	Q8	Q9	Q10
True	0.829	0.058	0.211	0.131	0.056	0.491	0.209	0.289	0.133

Table 1: Rate of true answers for annotation questions Q2–Q10 across the corpus.

	Q1	Q2**	Q3**	Q4**	Q5	Q6**	Q7**	Q8**	Q9*	Q10**
F1	0.74	0.86	0.76	0.85	0.88	0.92	0.64	0.83	0.77	0.92

Table 2: SVM classification accuracy: bag-of-words features; 80-20 train-test split; 5-fold cross validation. For the first question, this distinguishes highlighted text vs. its complement (excluded vs. included). For the rest of the questions, this distinguishes text of true instances from text of false instances, and is different from majority class baseline \* at  $p < 0.04$ ,  $t = -3.5$  and \*\* at  $p < 0.01$ ,  $t > |2.49|$ .

of cases. This suggests that subjects construe imagery as backgrounded or peripheral to the document, much as speakers regard co-speech iconic gesture as peripheral to speech (Schlenker and Chemla, 2017). Note, by contrast, that subjects characterized 12.7% of images as introducing a new tool: this includes many cases where the same subjects say the image is not required. In other words, subjects’ intuitions suggest that coherent imagery typically does not contribute instruction content, but rather serves as a visual signal that facilitates inferences that have to be made to carry out the instruction regardless. Our annotated examples, where imagery is linked to specific kinds of inferences, provide materials to test this idea.



TEXT: Top with another layer of ravioli and the remaining sauce not all the ravioli may be needed. Sprinkle with the Parmesan.

Figure 2: The image depicts both the action and the result of the action. The recipe is from Autodesk Inc. [www.instructables.com](http://www.instructables.com) and was contributed by [www.RealSimple.com](http://www.RealSimple.com).

**The Complex Coherence of Imagery.** Our annotation reveals cases where a single image does include more information than could be packaged into a single textual discourse unit (the proverbial thousand words). In particular, such imagery participates in more complex coherence relationships than we find between text segments. Multiple temporal relationships show this most clearly: 12% of

images that have any temporal relation have more than one. For example, many images depict the action that is described in the text, while also showing preparations that have already been made by displaying the scene in which the action is performed. Figure 2 depicts the action and the result of the action. It also shows how to prepare before carrying out the action. Other images show an action in progress but nearing completion and thereby depict the result. For instance, the image that accompanies “mix well until blended” can show both late-stage mixing and the blended result. Looking at a few such cases closely, the circumstances and composition of the photos seem staged to invite such overlapping inferences.

Such cases testify to the richness of multimodal discourse, and help to justify our research methodology. The True/False questions characterize the relevant features of interpretation without necessarily mapping to single discourse relations. For instance, Q4 and Q5 indicate inferences in line with an **Elaboration** relation; Q9 and Q10 indicate inferences in line with an **Exemplification** relation, as information presented in images show just one case of a generalization presented in accompanying text. However, our data shows that these inferences can be combined in productive ways, in keeping with the potentially complex relevant content of images.

**Information across modalities.** We carried out machine learning experiments to assess what information images provide and what textual cues can guide image interpretation. We use SVM classifiers for performance, and Multinomial Naive Bayes classifiers to explain classifier decision making, both with bag-of-words features.

Table 2 reports the F1 measure for instance classification with SVMs (with 5-fold cross validation). In many cases, machine learning is able to find cues that reliably help guess inferential pat-



Q4. Text has quantities not in image			
True		False	
1	-4.1	add	-4.5
cup	-4.4	place	-4.9
minutes	-4.7	put	-5.0
2	-4.7	make	-5.1
1/2	-4.9	mix	-5.1

Q8. Image depicts action in progress			
True		False	
add	-5.0	1	-4.6
mix	-5.2	cup	-4.7
place	-5.3	minutes	-4.9
bread	-5.5	160	-5.1
make	-5.6	put	-5.2

Table 3: Top five features of Multimodal Naive Bayes classifier for two classification problems and their corresponding log-probability estimates.

terns. Table 3 looks at two effective Naive Bayes classifiers, for Q4 (text has quantities) and Q8 (image depicts action in progress). It shows the features most correlated with the classification decision and their log probability estimates. For Q4, not surprisingly, numbers and units are positive instances.

More interestingly, verbs of movement and combination are negative instances, perhaps because such steps normally involve material that has already been measured. For Q8, a range of physical action verbs are associated with actions in progress; negative features correlate with steps involved in actions that don't require ongoing attention (e.g., baking). Table 4 reports top SVM with NB (NBSVM) (Wang and Manning, 2012) features for Q1 that asks subjects to highlight the part of the text that is most related to the image. Action verbs are part of highlighted text, whereas adverbs and quantitative information that cannot be easily depicted in images are part of the remaining segments of the text. Such correlations set a direction for designing or learning strategies to select when to include imagery.

## 5 Conclusions

In this paper, we have presented the first dataset describing discourse relations across text and imagery. This data affords theoretical insights into the connection between images and instructional text, and can be used to train classifiers to support automated discourse analysis. Another important

Q1. Information in text		
1	do it clearly	on which
2	let cool for	favorite toppings
3	recipe with directions	after an
4	how slowly the	lightly season
5	7 minutes on	the 2

Q1. Information in images		
1	added a beautiful	cover with
2	put as much	scrapping the
3	skin off of	finally fold
4	cut side toward	after an
5	blend and blend	add a

Table 4: Top five bigram and trigram features of NB-SVM for the first question. The highlighted text that is most relevant to the image describes depicted actions, while the complement descriptions describe quantities or modifications of the actions that are described in the highlighted segments.

contribution of this study is that it presents a discourse annotation scheme for cross-modal data, and establishes that annotations for this scheme can be procured from non-expert contributors via crowd-sourcing.

Our paper sets the agenda for a range of future research. One obvious example is to extend the approach to other genres of communication with other coherence relations, such as the distinctive coherence of images and caption text (Alikhani and Stone, 2019). Another is to link coherence relations to the structure of multimodal discourse. For example, our methods have not yet addressed whether image-text relations have the same kinds of subordinating or coordinating roles that comparable relations have in structuring text discourse (Asher and Lascarides, 2003). Ultimately, of course, we hope to leverage such corpora to build and apply better models of multimodal communication.

## Acknowledgments

The research presented here is supported by NSF Award IIS-1526723 and through a fellowship from the Rutgers Discovery Informatics Institute. Thanks to Gabriel Greenberg, Hristiyan Kourtev and the anonymous reviewers for helpful comments. We would also like to thank the Mechanical Turk annotators for their contributions.

## References

- Dorit Abusch. 2013. Applying discourse semantics and pragmatics to co-reference in picture sequences. In *Proceedings of Sinn und Bedeutung 17*, pages 9–25, Paris.
- M. Alikhani and M. Stone. 2018a. [Exploring coherence in visual explanations](#). In *2018 IEEE Conference on Multimedia Information Processing and Retrieval (MIPR)*, pages 272–277.
- Malihe Alikhani and Matthew Stone. 2018b. Arrows are the verbs of diagrams. In *Proceedings of the 27th International Conference on Computational Linguistics*, pages 3552–3563.
- Malihe Alikhani and Matthew Stone. 2019. “caption” as a coherence relation: Evidence and implications. In *Second Workshop on Shortcomings in Vision and Language (SiVL)*.
- Nicholas Asher and Alex Lascarides. 2003. *Logics of conversation*. Cambridge University Press.
- Mousumi Banerjee, Michelle Capozzoli, Laura McSweeney, and Debajyoti Sinha. 1999. Beyond kappa: A review of interrater agreement measures. *Canadian journal of statistics*, 27(1):3–23.
- Anne Cocos, Aaron Masino, Ting Qian, Ellie Pavlick, and Chris Callison-Burch. 2015. Effectively crowdsourcing radiology report annotations. In *Proceedings of the Sixth International Workshop on Health Text Mining and Information Analysis*, pages 109–114.
- Neil Cohn. 2013. Visual narrative structure. *Cognitive science*, 37(3):413–452.
- Heriberto Cuayáhuatl, Simon Keizer, and Oliver Lemon. 2015. Strategic dialogue management via deep reinforcement learning. In *NIPS Workshop on Deep Reinforcement Learning*. ArXiv:1511.08099.
- Samuel Cumming, Gabriel Greenberg, and Rory Kelly. 2017. Conventions of viewpoint coherence in film. *Philosophers’ Imprint*, 17(1):1–29.
- Steven K Feiner and Kathleen R McKeown. 1991. Automating the generation of coordinated multimedia explanations. *Computer*, 24(10):33–41.
- Joseph L Fleiss and Jacob Cohen. 1973. The equivalence of weighted kappa and the intraclass correlation coefficient as measures of reliability. *Educational and psychological measurement*, 33(3):613–619.
- Tuomo Hiippala and Serafina Orekhova. 2018. Enhancing the ai2 diagrams dataset using rhetorical structure theory.
- J. Hunter, N. Asher, and A. Lascarides. 2015. Integrating non-linguistic events into discourse structure. In *Proceedings of the 11th International Conference on Computational Semantics (IWCS)*, pages 184–194, London.
- Alex Lascarides and Matthew Stone. 2009. Discourse coherence and gesture interpretation. *Gesture*, 9(2):147–180.
- Scott McCloud. 1993. *Understanding comics: The invisible art*. William Morrow.
- Eleni Miltsakaki, Rashmi Prasad, Aravind K. Joshi, and Bonnie L. Webber. 2004. The Penn Discourse Treebank. In *LREC*. European Language Resources Association.
- Rashmi Prasad, Nikhil Dinesh, Alan Lee, Eleni Miltsakaki, Livio Robaldo, Aravind K. Joshi, and Bonnie L. Webber. 2008. The Penn Discourse TreeBank 2.0. In *LREC*. European Language Resources Association.
- Philippe Schlenker and Emmanuel Chemla. 2017. Gestural agreement. *Natural Language & Linguistic Theory*, pages 1–39.
- Matthew Stone and Una Stojnic. 2015. Meaning and demonstration. *Review of Philosophy and Psychology*, 6(1):69–97.
- Anthony J Viera, Joanne M Garrett, et al. 2005. Understanding interobserver agreement: the kappa statistic. *Fam Med*, 37(5):360–363.
- Sida Wang and Christopher D Manning. 2012. Baselines and bigrams: Simple, good sentiment and topic classification. In *Proceedings of the 50th Annual Meeting of the Association for Computational Linguistics: Short Papers-Volume 2*, pages 90–94. Association for Computational Linguistics.
- Bonnie Webber. 2009. Genre distinctions for discourse in the penn treebank. In *Proceedings of the Joint Conference of the 47th Annual Meeting of the ACL and the 4th International Joint Conference on Natural Language Processing of the AFNLP: Volume 2-Volume 2*, pages 674–682. Association for Computational Linguistics.
- Semih Yagcioglu, Aykut Erdem, Erkut Erdem, and Nazli Ikizler-Cinbis. 2018. RecipeQA: A challenge dataset for multimodal comprehension of cooking recipes. In *EMNLP*, pages 1358–1368. Association for Computational Linguistics.