

Revise and Resubmit: An Intertextual Model of Text-based Collaboration in Peer Review

Ilia Kuznetsov

UKP Lab

Technical University of Darmstadt

Department of Computer Science

<https://www.ukp.tu-darmstadt.de>

Jan Buchmann

Technical University of Darmstadt

Department of Computer Science

UKP Lab

Max Eichler

Technical University of Darmstadt

Department of Computer Science

UKP Lab

Iryna Gurevych

Technical University of Darmstadt

Department of Computer Science

UKP Lab

Peer review is a key component of the publishing process in most fields of science. Increasing submission rates put a strain on reviewing quality and efficiency, motivating the development of applications to support the reviewing and editorial work. While existing NLP studies focus on the analysis of individual texts, editorial assistance often requires modeling interactions between pairs of texts—yet general frameworks and datasets to support this scenario are missing. Relationships between texts are the core object of the intertextuality theory—a family of approaches in literary studies not yet operationalized in NLP. Inspired by prior theoretical work, we propose the first intertextual model of text-based collaboration, which encompasses three major phenomena that make up a full iteration of the review–revise–and–resubmit cycle: pragmatic tagging, linking, and long-document version alignment. While peer review is used across the fields of science and publication formats, existing datasets solely focus on conference-style

Action Editor: Wei Lu. Submission received: 31 May 2022; revised version received: 13 July 2022; accepted for publication: 21 July 2022.

<https://doi.org/10.1162/coli.a.00455>

© 2022 Association for Computational Linguistics

Published under a Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International (CC BY-NC-ND 4.0) license

*review in computer science. Addressing this, we instantiate our proposed model in the first annotated multidomain corpus in journal-style post-publication open peer review, and provide detailed insights into the practical aspects of intertextual annotation. Our resource is a major step toward multidomain, fine-grained applications of NLP in editorial support for peer review, and our intertextual framework paves the path for general-purpose modeling of text-based collaboration. We make our corpus, detailed annotation guidelines, and accompanying code publicly available.*¹

1. Introduction

Peer review is a key component of the publishing process in most fields of science: A work is evaluated by multiple independent referees—peers—who assess the methodological soundness and novelty of the manuscript and together with the editorial board decide whether the work corresponds to the quality standards of the field, or needs to be revised and resubmitted at a later point. As science accelerates and the number of submissions increases, many disciplines experience reviewing overload that exacerbates the existing weaknesses of peer review in terms of bias and reviewing efficiency. Past years have been marked by an increasing attention to the computational study of peer review, with NLP applications ranging from review score prediction (Kang et al. 2018) to argumentation analysis (Hua et al. 2019), and even first experiments in fully automatic generation of peer reviews (Yuan, Liu, and Neubig 2021).

We define **text-based collaboration** as a process in which multiple participants asynchronously work on a text by providing textual feedback and modifying text contents. Peer reviewing is a prime example of text-based collaboration: Similar to other fields of human activity, from lawmaking to business communication, paper² authors receive textual feedback on their document, process it, modify the paper text accordingly, and send it for another revision round. During this work, the participants need to address a range of questions. Is the feedback constructive, and what parts of the feedback text require most attention? What particular locations in the text does the feedback refer to? What has changed in the new version of the text, have the issues been addressed, and what parts of the text need further scrutiny? Answering these questions is hard, as it requires us to draw cross-document, **intertextual relations** between multiple, potentially long texts. Despite the great progress in finding single documents and extracting information from them, cross-document analysis is only starting to get traction in NLP (Caciularu et al. 2021). General frameworks and models to support text-based collaboration are yet to be established.

Treating text as an evolving entity created by the author and interpreted by the reader in the context of other texts is the core of the **intertextuality theory**—a family of approaches in literary and discourse studies not yet operationalized in NLP (Kristeva 1980; Broich 1985; Genette 1997; Steyer 2015). While the theoretical groundwork of the past decades provides a solid basis for such operationalization, the existing body of work is mostly dedicated to the literary domain and lacks terminological and methodological unity (Forstall and Scheirer 2019). Inspired by the theoretical work, in this article we propose a *joint intertextual model of text-based collaboration* that incorporates

¹ <https://github.com/UKPLab/f1000rd>.

² In this work we do not make distinctions between publication types; the terms “paper,” “article,” “manuscript, etc., are thereby used interchangeably.

three core phenomena covering one full document revision cycle: (1) **Pragmatic tagging** classifies the statements in text according to their communicative purpose; (2) **Linking** aims to discover fine-grained connections between a pair of texts; (3) **Version alignment** aims to align two revisions of the same text. Creating this model requires us to revisit the notion of text commonly accepted in NLP, and to adopt a new *graph-based data model* that reflects document structure and encompasses both textual and non-textual elements crucial for human interpretation of texts. Our proposal is coupled with an implementation that allows extending the proposed data model to new document formats and domains.

Peer review is used by most fields of science, and both reviewing standards and publishing practices show significant variation across research communities. For example, the temporal restrictions of conference peer review reduce the number of potential revisions a manuscript might undergo; the continuity of research journals, on the other hand, allows the manuscript to be reviewed, revised, and resubmitted multiple times before acceptance. Pre-publication review assumes that only accepted manuscripts get indexed and distributed; post-publication review happens after the publication, making accessible the papers that would otherwise be discarded. Finally, while most of the peer reviewing is closed and anonymized, some communities opt for open review, including the disclosure of the identities of authors and reviewers.

All these factors have substantial effects on the composition of the peer reviewing data, along with discipline-specific reviewing practices and quality criteria. However, existing NLP datasets of peer reviews (Kang et al. 2018; Hua et al. 2019; Ghosal et al. 2022, etc.) are exclusively based on conference-style, pre-publication review in machine learning. To address this gap, we introduce the *F1000Research Discourse corpus* (F1000RD)—the first multidomain corpus in journal-style, post-publication, open peer review based on the F1000Research platform.³ Based on this corpus, we instantiate our intertextual model in the peer reviewing domain and conduct annotation studies in pragmatic tagging, linking, and version alignment for peer reviews and research papers, producing a novel *multilayered dataset* and providing *key insights* into the annotation of intertextual phenomena. We finally apply our proposed framework to investigate the interaction between different types of intertextual relations. Our new resource is a major step toward multidomain applications of NLP for peer reviewing assistance, and an exemplary source to guide the development of general models of text-based collaboration in NLP.

In summary, this work contributes:

- A theoretically inspired intertextual model of *text-based collaboration*;
- A new graph-based *data model* based on an extended notion of text;
- Three novel *tasks*: pragmatic tagging, linking, and version alignment;
- A richly annotated *corpus* of multidomain journal-style peer reviews, papers, and paper revisions;
- Practical *insights* and analysis of intertextual phenomena that accompany text-based collaboration during peer review.

³ <https://f1000research.com>.

The rest of the article is organized as follows: Section 2 provides the necessary background in NLP for peer reviewing and text-based collaboration and introduces intertextuality theory and the dimensions of intertextuality that guide the development of our proposed framework. Section 3 discusses the notion of text, introduces a novel graph-based data model well-suited for intertextual analysis, and formally specifies our proposed intertextual framework. This framework is instantiated in a corpus-based study of intertextuality in peer review in Section 4, where we introduce the peer reviewing workflow of F1000Research, describe the F1000RD corpus, and provide detailed insights into our annotation studies and their results. Section 5 follows up with a more general discussion of future research directions. Section 6 concludes the article with final remarks.

2. Background

2.1 NLP and Peer Review

Multiple strands of NLP research aim to improve the efficiency and fairness of peer review. A long-standing line of work in reviewer matching aims to generate reviewer–paper assignments based on the reviewers’ expertise (Mimno and McCallum 2007; Anjum et al. 2019). Pioneering the use of NLP for peer review texts, Kang et al. (2018) introduce the tasks of review score and paper acceptance prediction, sparking a line of follow-up work (Ghosal et al. 2019; Dycke et al. 2021). To compare reviewing practices between different communities, Hua et al. (2019) define the task of argumentation mining for peer reviews and devise a model that they then use to compare the composition of peer review reports at several major artificial intelligence (AI) conferences. Recent work by Yuan, Liu, and Neubig (2021) and Wang et al. (2020) pioneer the field of automatic peer review generation.

While work on review score and acceptance prediction based on the whole review or paper text is abundant, the applications of NLP to assist the process of reviewing itself are few: The argumentation mining approach of Hua et al. (2019) and the aspect and sentiment annotation by Yuan, Liu, and Neubig (2021) can be used to quickly locate relevant passages of the review report (e.g., questions or requests); the author response alignment approach suggested by Cheng et al. (2020) can assist reviewers, authors, and editors in disentangling discussion threads during rebuttal. We are not aware of prior work in NLP for peer reviews that models pragmatic role of peer review statements, links peer reviews to their papers, and compares paper versions—although those operations form the very core of the reviewing process and could greatly benefit from automation.

Peer review as a general procedure is used in most fields of science, but the specific practices, domain and topic distribution, reviewing standards, and publication formats can vary significantly across research communities and venues. Despite the methodological abundance, from the data perspective existing work in NLP for peer reviews focuses on a narrow set of research communities in AI that make their reviewing data available via the OpenReview⁴ platform. We are not aware of any multidomain datasets of peer reviews that represent a diverse selection of research communities.

Finally, although documents are discussed and revised in most areas of human activity, from lawmaking to education, the complex communication that surrounds the

4 <https://openreview.net>.

Table 1

Sources of peer reviewing data in NLP to date, and F1000Research. arXiv catalog covers physics, mathematics, computer science, statistics, and others. F1000Research hosts publications from a wide range of domains, from meta-science to medical case studies (see Section 4.5).

	ICLR	NeurIPS	arXiv	F1000Research
license	unclear	unclear	varied	CC-BY/CC0
reviews	yes	yes	no	yes
revisions	yes	no	yes	yes
domains	CS/AI	CS/AI	multi	multi

creation of texts often remains hidden and scattered across multiple communication channels. Peer review is an excellent source for the study of text-based collaboration, as it involves multiple parties reviewing and revising complex documents as part of a standardized workflow, with a digital editorial system keeping track of their communication and the corresponding text updates. To be useful for the study of text-based collaboration, this data should be made available under a clear, open license, and should include peer review texts as well as paper revisions.

A brief analysis of the existing sources of reviewing and revision data in NLP reveals that none of them meet those requirements: As of April 2022 the ICLR content published via OpenReview is not associated with a license; although NeurIPS publishes peer reviews for accepted papers, their revision history is not available; and although arXiv provides pre-print revisions and is clearly licensed, it does not offer peer reviewing functionality. Table 1 summarizes the prior sources of peer reviewing data and compares them to F1000Research, a multidomain open reviewing platform that we introduce in Section 4.

2.2 NLP and Text-based Collaboration

Our work focuses on the three core aspects of text-based collaboration: pragmatics, linking, and revision. Pragmatic tagging aims to assign communicative purpose to text statements and is closely related to the work in automatic discourse segmentation for scientific papers (Teufel 2006; Lauscher, Glavaš, and Ponzetto 2018; Cohan et al. 2019b). Linking draws connections between a text and its commentary, and is related to work in citation analysis for scientific literature: Citation contextualization (Chandrasekaran et al. 2020) aims to identify spans in the cited text that a citation refers to, and citation purpose classification aims to determine why a given work is cited (Cohan et al. 2019a). Version alignment is related to the lines of research in analysis of Wikipedia (Yang et al. 2017; Daxenberger and Gurevych 2013) and student essay revisions (Afrin and Litman 2018; Zhang and Litman 2015).

For all three aspects, existing work tends to focus on narrow domains and builds upon domain-specific task formulations: Discourse segmentation schemata for research papers are not applicable to other text types; citation analysis is facilitated by the fact that research papers use explicit inline citation style, so the particular citing sentence is explicitly marked—which is clearly not the case for most document types; and Wikipedia revisions are fine-grained and only cover few edits at a time, while in a general case a document might undergo substantial change in between revisions. Domain-independent, general task models of pragmatic tagging, linking, and version alignment are yet to be established.

Downloaded from http://direct.mit.edu/col/article-pdf/48/4/949/2061780/col_a_00455.pdf by guest on 23 December 2023

Moreover, treating those tasks in an isolated manner prevents us from modeling the interdependencies between them. Yet those interdependencies might exist: A passage criticizing the text is more likely to refer to a particular text location, which, in turn, is more likely to be modified in a subsequent revision, and the nature of the modification would depend on the nature of the commentary. To facilitate the cross-task investigation of intertextual relationships, a *joint* framework that integrates different types of intertextual relations is necessary—however, most related work comprises individual projects and datasets that solely focus on segmentation, linking, or revision analysis.

Establishing a joint framework for the study of intertextual relations in text-based collaboration would require a systematic analysis of intertextual relations that can hold. One such systematization is offered by the intertextuality theory—a family of works in literary studies that investigates the relationships between texts and their role in text interpretation by the readers, and we briefly review it below.

2.3 Intertextuality

Any text is written and interpreted in the context of other texts. The term “intertextuality” was coined by Kristeva (1980) and has since been refined, transformed, and reinterpreted by subsequent work. Although the existence of intertextual relationships is universally accepted, there exists little consensus on the scope and nature of those relationships, and a single unified theory of intertextuality, as well as a universal terminological apparatus, are yet to be established (Forstall and Scheirer 2019). Based on the prior theoretical work, we distill a set of *dimensions* that allow us to systematize intertextual phenomena and related practical tasks, and form the requirements for our proposed intertextual framework.

Intertextual relations can be categorized into (1) **types**. A widely quoted typology by Genette (1997) outlines five core intertextuality types: *intertextuality*^C (homonymous term redefined by Genette) is the literal presence of one text in another text, for example, plagiarism; *paratextuality* is the relationship between the text and its surrounding material, for example, preface, title, or revision history; *metatextuality* holds between a text and a commentary to this text, for example, a book and its critique; *hypertextuality* is loosely defined as any relationship that unites two texts together and is not metatextuality; and *architextuality* is a relationship between a text and its abstract genre prototype, for example, what makes text a poem, a news article, or a peer review report.

Intertextual relations vary in terms of (2) **granularity**, both on the source and on the target text side. Steyer (2015) summarizes this relation in terms of four *referential patterns*: a part of a text referring to a part of another text (e.g., quotation), a part of a text referring to a whole other text (e.g., document-level citation), a whole text referring to a part of the other text (e.g., analysis of a particular scene in a novel), and a whole text referring to a whole other text (e.g., a preface to a book). We note that while the specific granularity of a text “part” is of secondary importance from a literary point of view, it matters in terms of both linguistic and computational modeling, and finer-grained distinction might be desirable for practical applications.

Intertextual relations vary in terms of their (3) **overttness**. On the source text side, the intertextual nature of a passage might be signaled *explicitly* (e.g., by quotation formatting or citation) or *implicitly* (e.g., by referring to another text without any overt markers [Broich 1985]). The explicitness of the marker might vary depending on the granularity or the target passage: A text might not be referred to at all (which would constitute allusion or plagiarism), referred to as a whole (e.g., most citations in engineering sciences), or referenced with page or paragraph indication (e.g., most citations

Table 2

Examples of intertextual relations by type (archi-, meta-, and paratextuality) and overtness: explicit or implicit. Example of direct reference: “the discussion on page 5, line 9.” Example of indirect reference: “the discussion of fine-tuning approaches (somewhere in the paper).”

type	explicit	implicit
archi	document structure, templates	genre standards, pragmatics
meta	hyperlinks, citations, direct reference	allusion, plagiarism, indirect reference
para	edit history, manual diff	description of changes

in humanities or references to legal and religious texts), up to the level of individual lines and statements (e.g., in the fine-grained discussion during peer review). We note that the type of the overt marker does not need to match the granularity of the target passage: A research paper might refer to a particular sentence in the source work, but only signal it by a document-level citation.

The three dimensions of intertextuality—type, granularity, and overtness—form the basis of our further discussion and cover a wide range of intertextual phenomena. Table 2 provides further examples of intertextual relations across domains and use cases, focusing on the three intertextuality types relevant for our work: architextuality (pragmatic tagging), metatextuality (linking), and paratextuality (version alignment). It both demonstrates the scope of phenomena a general-purpose intertextual framework could cover, and the connections between seemingly unrelated aspects of intertextuality.

2.4 Data Models

A data model specifies how a certain object is represented and defines the way its elements can be related and accessed. Data models differ in terms of their expressivity: Although a more expressive data model naturally retains more information, in research a less expressive model is often preferable as it poses fewer requirements on the underlying data, and can thereby represent more object types in a unified fashion. This motivates the use of *linear text* as a de facto data model in NLP: A research paper, a news post, and a Tweet can all be converted into a linear sequence of characters, which can be left as is and accessed by character offset (as in most modern datasets) or further segmented into tokens and sentences which then become part of the data model, as in most “classic” NLP datasets like Penn Treebank (Marcus, Marcinkiewicz, and Santorini 1993), Universal Dependencies corpora (Nivre et al. 2016), and others.

Yet, a closer look at the phenomena covered by the intertextuality theory reveals that heavily filtered linear text might not be the optimal data model for studying cross-document discourse. Lacking a standard mechanism to represent document structure, linear model of text is not well suited for modeling phenomena on varying granularity levels—yet humans readily use text structure when writing, reading, and talking about texts. The de facto data model does not offer standardized mechanisms to represent (or at least, preserve) non-textual information, like illustrations, tables and references—yet those often form an integral part of the document crucial for text interpretation and the surrounding communication. Finally, while it is possible to draw cross-document connections between arbitrary spans in plain text ad hoc, standardized mechanisms for representing cross-document relations are missing. All in all, while well-suited for representing grammatical and sentence-level phenomena, the current approach to text preprocessing at least invites a careful reconsideration.

Downloaded from http://direct.mit.edu/col/article-pdf/48/4/949/2061780/col_a_00455.pdf by guest on 23 December 2023

First steps in this direction can be found in recent work. In a closely related example, Lo et al. (2020) introduce the S2ORC data model that allows unified encoding of document metadata, structure, non-textual and bibliographic information for scientific publications. This information is then used in a range of studies in document-level representation learning (Cohan et al. 2020), citation purpose classification (Cohan et al. 2019a), and reading assistance (Head et al. 2021), demonstrating that non-linguistic elements and document structure are crucial for both devising better representations of text and assisting humans in text interpretation. The S2ORC data model is tailored to the idiosyncrasies of scientific publishing and publication formats. Our Intertextual Graph model introduced below is an attempt on a more general, inclusive model of text that offers a standardized way to represent document structure, encapsulate non-textual content, and capture cross-document relations—all while being applicable to a wide range of texts.

3. Proposed Framework

We now introduce our proposed intertextual framework for modeling text-based collaboration, which we later instantiate in our study of peer reviewing discourse. As our prior discussion shows, modeling text-based collaboration poses a range of requirements our framework should fulfil. Text-based collaboration is not specific to a particular domain and readily crosses domain and genre boundaries, so our framework needs to be (R1) general, that is, not tied to the particularities of domain-specific applications and data formats. As text-based collaboration involves multiple texts, our framework should be able to (R2) represent several texts simultaneously and draw intertextual relations between them. The framework should be suited for representing intertextual relations at (R3) different levels of granularity, and (R4) allow drawing relations between textual and non-textual content. Finally, the framework should enable (R5) joint representation of different intertextual phenomena to facilitate the study of dependencies between different types of intertextuality. Our further discussion of the proposed framework proceeds as follows: The data model defines the representation of input texts, which is then used to formulate task-specific models for pragmatic tagging (architextuality), linking (metatextuality), and version alignment (paratextuality).

3.1 Intertextual Graph (ITG)

The core of our proposed framework is the Intertextual Graph graph data model (ITG, Figure 1). Instead of treating text as a flat character or token sequence, ITG represents it as a set of nodes N^G and edges E^G constituting a graph G . Each node $n_i \in N^G$ corresponds to a logical element of the text. ITG supports both textual (e.g., paragraphs, sentences, section titles) and non-textual nodes (e.g., figures, tables and equations), allowing us to draw relationships between heterogeneous inputs. Nodes of the ITG are connected with typed directed edges $e \in E^G$. We denote an edge between nodes n_i and n_j as $e(n_i, n_j)$. We define three core edge types:

- next edges connect ITG nodes in the linear reading order, similar to the mainstream NLP data models discussed above;
- parent edges mirror the logical structure of the document and the hierarchical relationship between sections, subsections, paragraphs, and so forth;




Figure 1 Basic Intertextual Graph. Left: A full document ITG representing the logical structure of text on different levels via parent edges, with Levy and Goldberg (2014) as example document. Right top: Nodes can encapsulate textual, as well as non-textual information, like tables and figures. Right bottom: Three core edge types in ITG. link edges can be divided in further subtypes depending on the relationship, can connect nodes of different modality (text and table) and granularity (sentence and section), and can cross document boundaries.

- link edges represent additional, intertextual connections between graph nodes, ranging from explicit in-document references (e.g., to figure) to citations and implicit cross-document and version alignment links, as introduced below.

A node can have an arbitrary number of incoming and outgoing link edges, allowing the data model to represent many-to-many relationships; however, the next edges must form a connected list to represent the reading order, and the parent edges must form a tree to represent the logical structure of the text.

The proposed graph-based data model has multiple advantages over the traditional, linear representation of text widely used in NLP: It offers a standardized way to draw connections between multiple texts (R2); text hierarchy represented by the parent edges provides a natural scaffolding for modeling intertextual relations at different granularity levels (R3); a graph-based representation allows for encapsulating non-textual information while still making it available for intra- and intertextual reference (R4); and the framework can represent different intertextuality types jointly, as we demonstrate later (R5). Our data model is generally applicable to any text (R1) as it does not rely on domain-specific metadata (e.g., abstracts, keywords) or linking and referencing behaviors (e.g., citations); the next edge mechanism retains the linear reading order of text and makes the documents represented within our model compatible with NLP approaches tailored to linear text processing, for example, pre-trained encoders like BERT (Devlin et al. 2019), RoBERTa (Liu et al. 2019), Longformer (Beltagy, Peters, and Cohan 2020), and so on.

Data models in NLP differ in terms of the units of analysis they enable; those units are used to devise representations of target phenomena: For example, syntactic parsing operates with tokens within one sentence, while text classification operates with whole documents. The hierarchical representation offered by the ITG allows flexibility

in terms of the unit of analysis: In principle, an intra- or intertextual relationship can be drawn between a single token and a whole document (e.g., citations), or between a sentence and a table. While document is the largest unit of analysis in our data model, the definition of the smallest unit remains open and can be chosen on the application basis. Paragraph is the smallest unit of written text that does not require non-trivial preprocessing, and is an attractive first choice for the tasks where paragraph-level granularity is sufficient. However, most discourse analysis tasks, including pragmatic tagging and linking discussed below, require finer granularity, and in this work we use *sentence as a minimal unit of analysis*. We point that this is a practical choice, and not a limitation of the proposed framework per se: The internal contents of a sentence can be represented as a sequence of word nodes connected via the *next* edges—or as a character sequence referenced by offset, in accordance with the current mainstream approach to text representation.

3.2 Architextuality and Pragmatics

We now turn to pragmatic tagging as reflection of the text's architextual relationship to its genre. The original definition of the term by Genette (1997) is proposed in the context of literary studies and encompasses other genre-defining elements like style, form, and declarative assignment ("*A Poem*"); here, we focus on the discourse structure as reflection of the genre. Text genres emerge as written communication standards, and to enable efficient communication, a text should adhere to the discourse structure imposed by its genre: A news article seeks to quickly convey new information, motivating the use of the pyramid structure where the most important new facts are placed in the beginning of the text; a research article, on the other hand, aims to substantiate a claim for new knowledge, and is expected to clearly delineate related work and state its own contribution; an argumentative essay takes a stance on a claim and provides evidence that supports or refutes this claim, and so on. Being able to determine the pragmatic structure of a text is the first key step to its interpretation.

Unlike the other two relationships described below, architextuality holds not between a pair of texts, but between a text and its abstract "prototype." To reflect this, we introduce the task of pragmatic tagging, where an ITG node n can be associated with a label from a pre-defined set $label(n) = l_i, \{l_1, l_2 \dots l_j\} \in L$. Pragmatic structure of text can be signalled by the overt text structure, for example, standardized section headings or text templates—in which case the architextual relationship to the text genre is explicit; more frequently, however, it remains implicit and needs to be deduced by the reader. A pragmatic label can be associated with the units of varying granularity, from a text section (e.g., "*Related work <...>*" → Background) to a particular sentence (e.g., "*The paper is well written and the discussion is solid.*" → Strength).

Pragmatic tagging is a generalization of a wide range of discourse tagging tasks, including argumentative zoning (Teufel 2006) and subtasks of argumentation mining (Stab and Gurevych 2014, 2017; Habernal and Gurevych 2017), and is related to work in rhetorical structure analysis (Mann and Thompson 1988). We note that our definition of pragmatic tagging does not cover the structural relationships between segments of the *same* text (like argumentative structures in argumentation theory or rhetorical relations in rhetorical structure theory)—while drawing such relations is easy within the proposed data model, this goes beyond the scope of our analysis. We instantiate pragmatic tagging in the peer reviewing domain by introducing a novel labeling schema for peer review analysis in Section 4.4.

3.3 Metatextuality and Linking

The next intertextuality type that our proposed framework incorporates is metatextuality, defined by Genette as a “*commentary. . . uniting a given text to another, of which it speaks without necessarily citing it.*” Despite its original use in the discussion of literary works, metatextuality lies at the very core of text-based collaboration, and spans beyond the literary domain: A book review, a related work survey, a social network commentary, or a forum thread post all participate in a metatextual relationship. Being able to draw such relationships between two texts is crucial for text interpretation—however, as metatextual relationships are not always explicitly signalled, this might often present a challenge that can greatly benefit from analysis and automation.

To model metatextuality in our framework, we introduce the task of linking. Given two ITGs, the anchor graph G^A and the target graph G^T , we use superscript notation to distinguish their nodes (e.g., $n_i^A \in N^A$). The goal of linking is then to identify the anchor nodes in G^A and draw link-type edges between their corresponding target nodes in G^T , $e(n_i^A, n_j^T)$. Linking is a frequent phenomenon, and while some text genres enforce explicit linking behavior (e.g., citations in scientific literature), in most texts the linking is done implicitly (e.g., mentioning the contents of the target text). Contrary to Genette’s definition, our interpretation of explicit linking subsumes the cases of direct text reuse via quotation. Links can vary greatly in terms of both source and target granularity: A sentence might link to a whole text or a particular statement in this text; and a paragraph of the anchor text might be dedicated to a single term mentioned in the target text. Links are frequently drawn between textual and non-textual content: For example, a sentence might refer to a table, and a social media post might comment on a video. Although our work does not deal with multimodality, the encapsulation offered by the ITG data model enables such scenarios in the future.

The task of linking is a direct generalization of a wide spectrum of existing NLP tasks covering specific use cases and narrow domains: Citation intent prediction (Cohan et al. 2019a; Teufel, Siddharthan, and Tidhar 2006), citation contextualization (Chandrasekaran et al. 2020), argument pair extraction (Cheng et al. 2020), review–paper mapping (Wang et al. 2020), and pairing comments on Wikipedia talk pages with the corresponding revision edits (Daxenberger and Gurevych 2014) can all be cast as link-finding or link-labeling within our proposed framework. Linking relates to a broad category of evidence-finding NLP tasks like question answering with span prediction (Kwiatkowski et al. 2019; Dasigi et al. 2021) or fact checking with evidence retrieval (Thorne et al. 2018; Wadden et al. 2020). Conceptually, linking is related to the work in cross-document structure theory (CST) (Radev 2000; Radev, Otterbacher, and Zhang 2004; Maziero, del Rosario Castro Jorge, and Salgueiro Pardo 2010), which has been applied to short documents on sentence level in the newswire domain. While CST focuses on *labeling* cross-document links and devising a typology of cross-document relations, in this work we focus on *finding* links between documents, and instantiate implicit and explicit linking in the peer reviewing domain in Section 4.5.

3.4 Paratextuality and Version Alignment

The final intertextuality type discussed here is paratextuality. Genette (1997) broadly defines paratextuality as a relationship between the core text of a literary work and the elements that surround it and influence the interpretation, including title, foreword, notes, illustrations, book covers, and so forth. We focus on a particular paratextual relationship highly relevant for modeling text-based collaboration—the relationship

between a text and its previous versions. An updated text is not interpreted anew, but in the context of its earlier version; being able to align the two is critical for efficient editorial work, as it would allow quick summarization of the changes and highlighting of the new material. Those time-consuming operations are mostly performed manually, as general-purpose models of text change are missing.

To address this, we introduce the task of version alignment. Given two ITGs corresponding to the different versions of the *same* text G^t and $G^{t+\Delta}$, the goal is to produce an alignment, which we model as a set of intertextual edges $e(n_i^{t+\Delta}, n_k^t)$ between the two graphs.⁵ Note that in our formulation the two versions of the document must not be consecutive, and the ability of the ITG to represent multiple documents allows us to simultaneously operate with multiple versions of the same document with an arbitrary number of revisions in-between. We further denote the revisions as *short-scope* or *long-scope*, depending on the magnitude of changes between them. Although not a strict definition, a typo correction constitutes a short-scope edit, whereas a major rewrite constitutes a long-scope edit.

In terms of overttness, the correspondence between two versions of the same text is rarely explicit: Producing such alignment manually is time-consuming, and the logs that keep track of character-level edit operations are limited to few collaborative authoring platforms like Google Docs⁶ and Overleaf,⁷ and are too fine-grained to be used directly. The alignment between two text versions thereby remains implicit: While the general fact of text change is known, the exact delta needs to be discovered via ad-hoc application of generic utilities like `diff` or by soliciting textual summaries of changes from the authors. Those might differ in terms of granularity from high-level notes (“*We have updated the text to address reviewer’s comments*” to in-detail change logs (“*Fixed typos in Section 3; adjusted the first paragraph of the Introduction.*”)); the choice of the granularity level depends on the application and the communicative scenario.

Version alignment is related to multiple strands of research within and outside NLP. Outside NLP, version analysis is explored in the software engineering domain (Moreno et al. 2017; Jiang, Armaly, and McMillan 2017)—which focuses on program code; related approaches based on simple text matching techniques exist in digital humanities, termed as *collation* (Nury and Spadini 2020). In NLP, Wikipedia edits and student essay writing have been the two prime targets for the study of document change. Both existing lines of research operate under narrow domain-specific assumptions about the nature of changes: Wikipedia-based studies (Yang et al. 2017; Daxenberger and Gurevych 2014) assume short-scope revisions characteristic of collaborative online encyclopediae, and focus on edit classification, whereas essay analysis (Afrin and Litman 2018; Zhang and Litman 2015) focuses on the narrow case of student writing and medium-sized documents. Our task definition generalizes from those previously unrelated strands of research and allows the study of long-scope long-document revisions, instantiated in the annotation study of research paper alignment in Section 4.6.

3.5 Joint Modeling

Apart from suggesting general, application-independent architectures for pragmatic tagging, linking, and version alignment of arbitrary texts, our framework allows joint

⁵ The direction of the version alignment is opposite to time (i.e., from $t + \Delta$ to t): A latter text refers to an earlier text, but not vice versa.

⁶ <https://docs.google.com>.

⁷ <https://overleaf.com>.

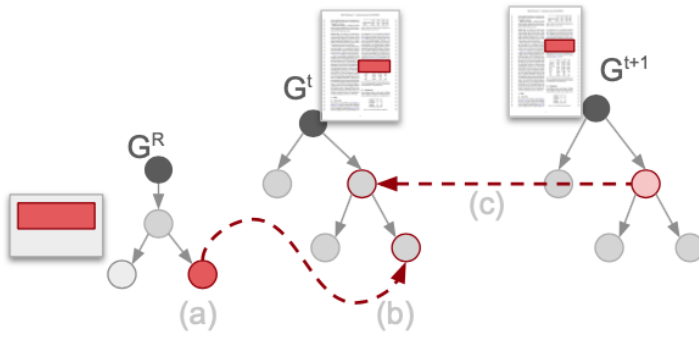


Figure 2 Joint modeling of multiple documents via ITG, simplified and with next edges omitted for clarity. A review G^R node with certain pragmatics (a, pragmatic tagging) is connected (b, linking) to the main document G^t which is later revised, producing a new version G^{t+1} aligned to the original (c, version alignment). Following the links between documents enables new types of intertextual analysis—yet the links are not always explicit (see Section 2.3), and might need to be inferred from text.

modeling of these phenomena (Figure 2). Different types of intertextuality indeed interact: The communicative scenario that a text serves does not only prescribe its pragmatic structure, but also determines the standards of linking and the nature of updates a text might undergo. On a finer level, joint modeling of pragmatics, linking, and version alignment allows us to pose a range of new research questions. Are metatextual statements with certain pragmatics more likely to be linked, and do statements with a large number of links tend to belong to a certain pragmatic category? Can explicit, readily available intertextual signals—document headings, citations, and detailed, character-level change logs—be used as auxiliary signals for uncovering latent, implicit intertextual relationships? What parts of texts are more likely to be revised, and which factors contribute to this? Our proposed framework facilitates joint analysis of intertextuality outside of narrow application-driven scenarios like using research paper structure to boost citation intent classification performance in Cohan et al. (2019a). We demonstrate this capacity in our study of peer reviewing discourse, which we now turn to.

4. Corpus Study in Peer Review

Research publication is the main mode of communication in the modern scientific community. As more countries, organizations, and individuals become involved in research, the number of publications grows, and although significant progress has been achieved in terms of finding research publications (Cohan et al. 2020; Esteva et al. 2021) and extracting information from them (Luan et al. 2018; Nye et al. 2020; Kardas et al. 2020), technologies for prioritizing research results and ensuring the quality of research publications are lacking. The latter two tasks form the core of peer review—a distributed manual procedure where a research publication is evaluated by multiple independent referees who assess its soundness, novelty, readability, and potential impact.

As a result of peer reviewing, the manuscript is accepted and published, or rejected and discarded, or revised and resubmitted for another reviewing round. During this process, the authors, reviewers, and volume editors work together to improve the manuscript so that it adheres to the scientific and publishing standards of the field.

The communication happens over text, making peer review a prime example of text-based collaboration. The process often takes place in a centralized digital publishing platform that stores a full log of the interactions between the participants, including initial manuscript draft, peer reviews, amendment notes, meta-reviews, revisions, and author responses; this makes peer review a unique, rich data source for the study of intertextual relations.

Due to the anonymity of the process, however, this data often remains hidden. As discussed in Section 2.1, the few existing sources of peer reviewing data used in NLP are insufficient to support the intertextual study of peer reviewing due to the gaps in domain and data type coverage and the lack of clear licensing. To approach this, we introduce F1000Research as a new peer reviewing data source for NLP. While meeting all of our requirements, the F1000Research platform has other substantial differences from the reviewing and publishing platforms previously used in NLP research on peer reviews. We briefly outline the reviewing process of F1000Research and highlight those differences below.

4.1 Data Source: F1000Research

F1000Research is a multidomain open access journal with fully open post-publication reviewing workflow. It caters to a wide range of research communities, from medicine to agriculture to R package development. Unlike regular, “closed” conferences and journals, F1000Research publishes the manuscripts directly upon submission, at which point they receive a DOI and become citable. After this, the authors or the F1000Research staff invite reviewers, who provide review reports and *recommendations* to approve, approve-with-reservations, or reject the submission.

Reviewers are presented with guidelines and domain-specific questionnaires, but reviews themselves are in a free-text format. Authors can write individual author responses and upload a new version that is reviewed by the same referees, producing a new round of reports. This “revision cycle” can repeat until the paper is approved by all referees. However, the official “acceptance decision” step common to traditional journals and conferences is not required here: A paper might be rejected by its reviewers and still be available and citable. Crucially for our task, the reviewing process at F1000Research is fully transparent, reviewer and author identities are public, and reviews are freely accessible next to the manuscript under an explicit CC-BY or CC-0 license. All the articles and reviews at F1000Research are available as PDF and as easy-to-process JATS XML,⁸ which allows us to avoid the noise introduced by PDF-to-text conversion and makes fine-grained NLP processing possible.

All in all, F1000Research provides a unique source of fine-grained peer reviewing data so far overlooked by the NLP community. In this work we focus on papers, paper revisions, and reviews from F1000Research, and leave a thorough exploration of author responses and revision notes to future work.

4.2 Corpus Overview

The full F1000RD corpus published with this work was crawled on April 22, 2021, from the F1000Research platform using the official API.⁹ Source JATS XML files were converted into the ITG representation as described in Section 3. We have collected peer

⁸ <https://jats.nlm.nih.gov>.

⁹ <https://f1000research.com/developers>.

Table 3

Paper first-version statistics for F1000RD full corpus vs. sample, as well as the number of papers that have more than zero reviews for the first version, and more than one version. Number of words here and further is the lower bound estimated via whitespace tokenization.

	papers	#words	#sentences	+reviews	+revisions
full	5.4k	17.4M	–	3.7k	1.6k
sample	172	496K	24.2K	172	122

review texts, papers, and revisions for each paper available at F1000Research at the time of the crawl. The resulting full dataset contains 5.4k papers, of which 3.7k have reviews and 1.6k have more than one version (Table 3). This makes our resource comparable to the widely known PeerRead dataset (Kang et al. 2018), which contains approximately 3k papers with reviews. Table 5 provides basic statistics for the peer review part of the dataset; the number of reviews is slightly lower but comparable to PeerRead (10k). We note the high proportion of accepted papers in both datasets; however, the accept-with-reservations mechanic is specific to F1000Research and allows us to collect more critical reviews that contain actionable feedback.

We have selected a sample from the full dataset for the in-detail investigations described in the following sections. To avoid domain bias, we would like our sample to include contributions from different disciplines and publication types; experiments in pragmatic tagging and linking require publications that have at least one peer review for the first version of the manuscript; experiments in version alignment additionally require the manuscript to have at least one revision. While the latter criteria are easily met via filtering, F1000Research does not enforce an explicit taxonomy of research domains. Instead, F1000Research operates with **gateways**—collections of publications that loosely belong to the same research community.¹⁰ To ensure a versatile multidomain and multiformat sample, for this work we have selected publications from the following gateways. While it is possible for a publication to belong to multiple gateways, we have only selected publications assigned to a single gateway.

- Science policy research (*scip*) publishes manuscripts related to the problems of meta-science, peer review, incentives in academia, etc.
- ISCB Community Journal (*iscb*) is the outlet of the International Society for Computational Biology dedicated to bioinformatics and computational biology.
- RPackage (*rpkg*) publishes new R software package descriptions and documentation; those undergo community-based journal-style peer review and versioning.
- Disease outbreaks (*diso*) contains research articles in the public health domain; many recent publications are related to the COVID pandemic, vaccination programs, and public response.

¹⁰ <https://f1000research.com/gateways>.

Table 4
Manuscript statistics in the F1000RD sample by domain.

	case	diso	iscb	scip	rpkg	total
papers	45	37	31	31	28	172
#words	63K	118K	111K	117K	88K	496K

Table 5
Review statistics for base F1000RD full corpus vs. sample, with ratios for approve, approve-with-reservations, and rejecting reviews.

	reviews	approve	approve-w-r	reject	#words	#sentences
full	8,053	.55	.38	.07	2M	–
sample	224	.36	.53	.11	59K	4.9K

- Medical case reports (*case*) are a special publication type at F1000Research, but do not constitute a separate gateway; they mostly describe a single clinical case or patient, often focusing on rare conditions or new methodology and treatment.

Tables 3 and 5 compare our sample to the full corpus. As evident from Table 5, the study sample contains a lower proportion of straight approve reviews, focusing on the reports that are more likely to discuss the submission in detail, link to the particular locations in the submission, and trigger a revision. Table 4 compares the gateways' contributions to our study sample. As it shows, the sample contains similar amounts of text for each of the gateways; the divergence among average manuscript length reflects the publication type differences (e.g. medical case reports contain 1.4k words on average, while scientific policy articles span an average of 3.7k words).

4.3 Preprocessing and Annotation Setup

In addition to converting source documents from JATS XML into the ITGs, the *reviews* in the F1000RD study sample were manually split into sentences; similar to Thelwall et al. (2020), we clean up review texts by removing template reviewing questionnaires included by the F1000Research reviewing interface. Because manually splitting *papers* into sentences would be too labor-intensive, for papers we used the automatic parses produced by *scispaCy*.¹¹

We add three layers of intertextual annotation to this data, illustrated in Figure 3, which instantiates our general-purpose model from Figure 2 in the peer reviewing domain. Pragmatic analysis of the peer reviews (a) allows us to determine the communicative purpose of individual reviewer statements. Linking between a peer review and a paper (b) allows us to find locations in the manuscript that a reviewers' commentary refers to (e.g., a paragraph dedicated to the experiment timing). Version alignment (c) allows us to trace the changes to these locations in the updated version (e.g., new

¹¹ <https://allenai.github.io/scispaCy>.

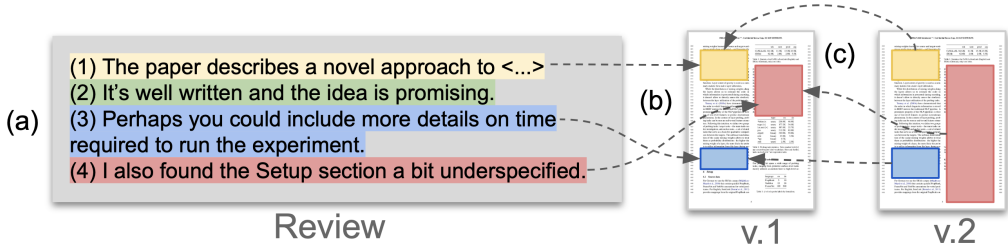


Figure 3 Annotation layers at a glance. For each sentence in the Review, we determine its pragmatics (a), e.g., Recap (1), Strength (2), Todo (3), or Weakness (4). We then link (b) the review sentences to the first version of the paper. Once the paper is updated, we can (c) align the first and second version to study the relationship between the review, the paper, and its revision.

details on the time and resources required to run the experiment). Our joint approach to modeling enables new types of analysis (e.g., establishing whether the reviewers’ feedback has triggered changes, or whether new, previously not reviewed content has been added to the revision).

Most annotations reported in this work were performed by two main annotators, both fluent non-native English speakers pursuing a Masters degree and comfortable with reading academic texts. We aimed to recruit annotators with diverse research backgrounds: One annotator had a background in Environmental Engineering, another one has previously studied Business Administration and was pursuing a degree in Data and Discourse Studies. When discussing the annotation studies, the further text refers to these two annotators, unless explicitly specified. Additional annotations were performed by the authors of this work, fluent non-native English speakers with extensive expertise in academic reading, as well as authoring and receiving peer reviewing feedback, with background in theoretical linguistics, computer science, as well as computational and molecular biology. To study the effect of domain expertise on linking (Section 4.5), we additionally involved two medical experts who had both completed their studies and had a full practicing year by the time of the annotation study.

The differences between pragmatic tagging, linking, and version alignment motivate the different approaches we took to annotate those phenomena. Pragmatic tagging was cast as a sentence labeling task, annotated by the two main annotators supplied with a guideline, and adjudicated by an expert. The annotation of implicit links was assisted by a suggestion module to reduce the number of potential linking candidates and lower the cognitive load on the main annotators who had to simultaneously handle two long documents during the linking process. The annotation of version alignment was performed automatically and later manually verified by the expert annotators. The following sections describe our annotation process in detail. The detailed annotation guidelines used for pragmatic tagging and linking, along with the resulting annotated data and auxiliary code, are available at <https://github.com/UKPLab/f1000rd>.

4.4 Pragmatic Tagging

While peer review is used in virtually every field of research, peer reviewing practices vary widely depending on the discipline, venue type, as well as individual reviewers’ experience level, background, and personal preferences. Figure 4 lists a range of reviews from ICLR-2019 (hosted by OpenReview) and F1000Research. As we can see, even

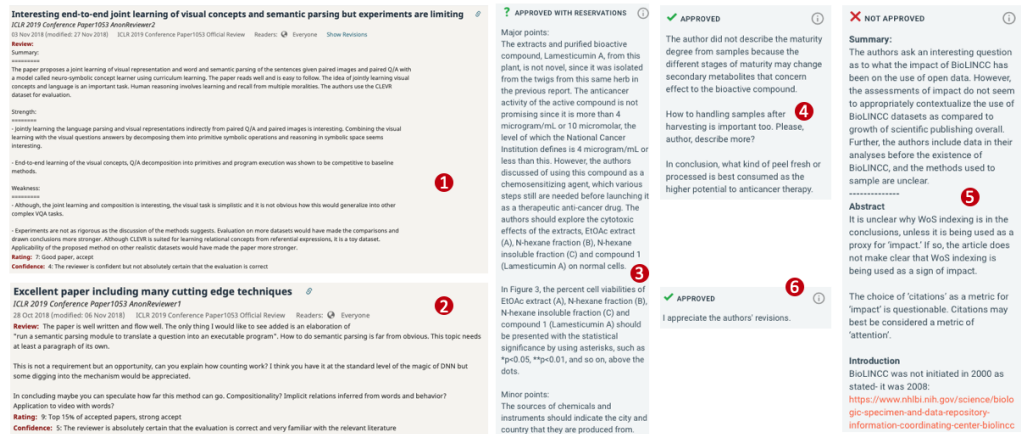


Figure 4 Diversity of peer reviewing styles and review report structures. Ex. 1 and 2: ICLR-2019 via OpenReview; Ex. 3–6: F1000Research.

within one venue reviews can differ dramatically in terms of length, level of detail, and writing style: Whereas reports 1, 3, and 5 are structured, reports 2, 4, and 6 are free-form; moreover, among the structured reports, the reviewer in 1 groups their comments into summary, strengths, and weaknesses; while the reviewer in 3 organizes their notes by priority (major and minor points); and the reviewer in 5 comments by article section. This illustrates the great variability of texts that serve as academic peer reviews.

Despite this variability, all peer reviewing reports pursue the same communicative purpose: to help the editor decide on the merit of the publication, to justify the reviewers' opinion, and to provide the authors with useful feedback. Uncovering the latent discourse structure of free-form peer reviewing reports has several applications: it might help control the quality of reviewing reports by detecting outliers (e.g., reports that mention no strengths or do not provide a paper summary) before or after the report is submitted; it might help editors and authors navigate the peer review texts and summarize feedback; or it can be used to compare reviewing styles across disciplines and venues similar to the argumentation analysis by Hua et al. (2019).

4.4.1 Task and Annotation. Motivated by this, we instantiate the task of pragmatic tagging $label(n) = l_i \in L$ in the peer reviewing domain with a sentence-level pragmatic tagging schema inspired by the related work in argumentation mining (Hua et al. 2019) and sentiment analysis (Yuan, Liu, and Neubig 2021; Ghosal et al. 2019) for peer reviews, as well as by the commonplace requirements from peer reviewing guidelines and templates. Our proposed six-class schema covers the major communicative goals of a peer reviewing report. Recap sentences summarize content of the paper, study, or resource without evaluating it; this includes both general summary statements and direct references to the paper. Weakness and Strength express an explicit negative or positive opinion about the study or the paper. Todo sentences contain recommendations and questions, something that explicitly requires reaction from the authors. Structure is used to label headers and other elements added by the reviewer to structure the text. Finally, an open class *Other* is used to label everything else: own reasoning, commentary on other publications, citations. Table 6 provides examples for each of the classes

Table 6
Pragmatic tagging schema, examples, and correspondence with the previously proposed AMPERE schema.

label	example	AMPERE
Recap	<i>The authors describe a case of <...></i>	Fact
Weakness	<i>The figures are of low quality.</i>	Evaluation
Strength	<i>It is a well written software article.</i>	Evaluation
Todo	<i>Please specify whether <...></i>	Request
Structure	<i>My major concerns:</i>	Other
Other	<i>As a non-surgeon, I can not <...></i>	Other, Fact, Evaluation

Table 7
Inter-annotator agreement (Krippendorff’s α) for pragmatic tagging by domain.

all	case	diso	iscb	rpkg	scip
0.77	0.78	0.75	0.77	0.74	0.79

and compares them to AMPERE (Hua et al. 2019): Whereas AMPERE focuses on surface-level argumentative structure, pragmatic analysis requires us to draw a distinction between Strengths and Weaknesses (both *Evaluation* in AMPERE) and to separate the discussion of the background from the discussion of the manuscript under review (both *Fact*).

To evaluate the robustness and coverage of our schema and produce the pragmatic tagging layer of the F1000RD corpus, we conducted an annotation study. After four hours of training, two annotators labeled the F1000RD study sample according to the schema. While regular structured discussion meetings have been scheduled throughout the annotation process, the labeling itself was done independently by the two annotators, who reached a substantial (Landis and Koch 1977) inter-annotator agreement of 0.77 Krippendorff’s α , demonstrating the high robustness of the proposed schema. Table 7 outlines the inter-annotator agreement for pragmatic tagging by domain; as we can see, despite the domain differences, the schema remains robust across domains, suggesting that our proposed pragmatic categories of peer reviewing can be reliably detected and labeled universally. The initial annotations were adjudicated by an expert annotator (author of the schema), who resolved disagreements between annotators and in rare cases harmonized the annotation to fit the final guidelines, taking into account the refinements made over the period of the annotation study.

4.4.2 Analysis. The resulting 4.9K labeled sentences make up the pragmatics layer of the F1000RD dataset. Figure 5 shows the distribution of pragmatic classes in the harmonized corpus. As we can see, the proposed annotation schema covers more than 82% of the review full text, with only 17% of the sentences falling into the catch-all *Other* category. Turning to the distribution of the core pragmatic classes, we note a high proportion of *Todo* sentences, whereas the related *Request* category in the AMPERE dataset makes up for less than 20% of the data (Hua et al. 2019, Table 3). One possible explanation for this discrepancy could be the difference between reviewing workflows: While AMPERE corpus builds upon conference-style reviewing data from ICLR with few revision rounds and quick turnaround, our data is based on the journal-style,




Figure 5
Distribution of pragmatic tags in F1000RD peer review annotations.

post-publication reviewing process of F1000Review, which allows the reviewers to make more suggestions for the next iteration of the manuscript. This finding highlights the importance of data diversity and calls for further comparative research in peer reviewing across domains, disciplines, and publishing workflows.

Due to the robustness of the schema, the two annotators agreed on the label for more than 80% of the sentences in the corpus, creating a large catalog of clear examples of sentences for each pragmatic class. Table 8 lists examples of clear cases by class; as it shows, the categories are natural enough to capture pragmatic similarities between sentences, while at the same time allowing for meaningful variation depending on the aspect in focus and the research domain. We note a difference in specificity and granularity of the sentences within the same pragmatic class, ranging from general statements (“*The paper is hard to read*”) to surgical, low-level edit suggestions (“*Implementation paragraph, line 6: pm-signature -> pmsignature.*”).

Table 8

Clear-case examples of pragmatic classes from the corpus; some sentences are shortened and the Other, Recap, and Structure classes are omitted for the sake of presentation.

Strength	<p>XLA is a rare disease which gave this case report high value for being indexed. It is a well written case report and the discussion is precise. Each step is clearly explained and is accompanied by R code and output (...) I appreciated the author explaining the role that preprints could play (...) Recycling the water in the traps is a good idea in the short term because (...)</p>
Weakness	<p>1.5 year follow up is short for taste disorders. This doesn't seem efficient, especially with large single cell data sets such as (...) The use of English is somewhat idiosyncratic and requires minor review. The conclusions, while correct, are weak, and the results section is missing. The way they support their claim is problematic.</p>
Todo	<p>How were Family vs. Domain types handled from InterPro or Pfam? I recommend to the author to delete in the discussion the sentence (...) The following important reference is missing, please include: (...) The role of incentives in CHW performance should be discussed: (...)</p>

The common sources of disagreement between annotators highlight the limitations of the proposed schema and point at the directions for future improvement. We observe a high number of disagreements between the Recap and Other categories due to the failure to distinguish between the manuscript itself, the accompanying materials, the underlying study, and background knowledge; this is especially pronounced in the medical and software engineering domains, which make frequent references to patient treatment and accompanying program code, respectively. Similarly, we observe a large number of disagreements stemming from the interpretation of a sentence as neutral or evaluating: For example, “*The authors have detailed the pattern seen in each case*” can be interpreted as a neutral statement describing the manuscript, or as a Strength. An important subset of the Other category not accounted for by our proposed schema is performative and includes meta-statements like “*I recommend this paper to be indexed*” and “*I don’t have sufficient expertise to evaluate this manuscript.*” We note that such statements align well with the common elements of structured peer reviewing forms—overall and confidence score—highlighting the connections between explicit and implicit dimensions of the peer reviewing pragmatics.

4.5 Linking

Metatextuality is deeply embedded in academic discourse: Each new study builds upon vast previous work, and academic texts are abundant with references to corresponding publications; the number of incoming references accumulated over time serves as a proxy of the publications’ influence, and the total reference count is a common measure of individual researchers’ success. The main mechanism of intertextual referencing in academic writing is the citation; while the core function of connecting a text to a related previous text is common across research disciplines, the specific citation practices vary among communities, from document-level citations in textbooks to precise, page- and paragraph-level inline references. Automatic analysis of citation behavior is a vast field of research in NLP (Cohan et al. 2019a; Teufel, Siddharthan, and Tidhar 2006; Chandrasekaran et al. 2020).

Like academic publications, peer reviews are also deeply connected to the manuscripts on an intertextual level. However, compared with full papers, peer reviews represent a much finer level of analysis as their main goal is to scrutinize a single research publication; moreover, since deep knowledge of the text is implied both from the author and from the reviewer side, most intertextual connections between the two texts remain implicit. Uncovering those connections bears great potential for automation and reviewing assistance, yet NLP datasets and methods to support this task are lacking.

4.5.1 Task. We instantiate the task of linking in the peer reviewing domain as follows: Given the ITGs of the peer review G^R and the paper G^P , our goal is to discover the intertextual links $e(n_i^R, n_j^P)$ between the anchor nodes in the review $n_i^R \in G^R$ and the target nodes in the paper $n_j^P \in G^P$. We distinguish between explicit and implicit linking and model them as two separate subtasks. The anchor of an explicit link contains an overt marker pointing to a particular element of the paper text. It can also contain a clearly marked quotation from the paper, for example, “*In the Introduction, you state that <...>*”. The anchor of an implicit link refers to the paper without specifying a particular location, for example, “*The idea to set up mosquito traps is interesting,*” and is a substantially more challenging task. Similar to pragmatic tagging, we take sentence as a minimum unit for the anchor nodes n_i^R ; the granularity of the target nodes n_j^P is

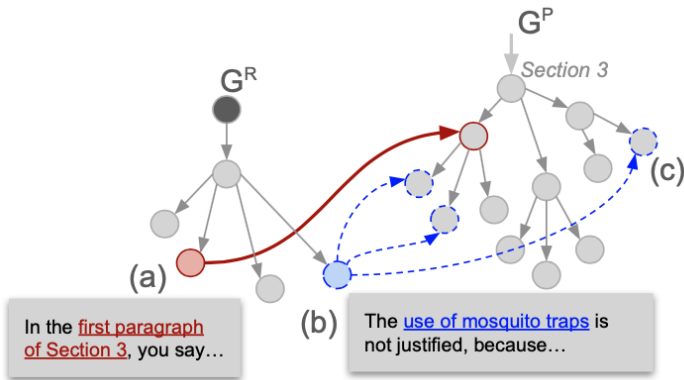


Figure 6

Linking between a review G^R and a paper G^P . Only one paper section shown for simplicity, next edges omitted. While explicit linking (a) is facilitated by the presence of a clear anchor (“first paragraph”) and clearly defined target in G^R , implicit linking faces two challenges at once, as it is unclear both whether the anchor sentence links to the paper at all (b), and if yes, what passages of the paper are linked to (c). Answering those questions requires simultaneous work with both documents; we use a suggestion-based approach to make annotation feasible.

Table 9

Examples for explicit and implicit links. Explicit anchors are underlined.

Review Anchor	Link Type	Paper Target (Node Type)
The most important part of the article is in the <u>discussion</u> .	explicit	Discussion (section)
Fig. 4 and <u>Table 3</u> , interpretation is not helped by the lack of correspondence between names and code {...}	explicit	Figure 4 (figure)
	explicit	Table 3 (table)
It would be good to have a set of images from CellProfiler.	implicit	Nuclei and infected cells were counted using CellProfiler. (sentence)
The authors intended to design a code requiring little R expertise.	implicit	We intentionally used simple syntax such that users with a beginner level of experience with R can adapt the code as needed. (sentence)
Details about the SVM learning algorithm must be included in the <u>methods</u> .	explicit	Methods (section)
	implicit	SVM learning: Previously, paclitaxel-related response genes were identified, and their expression in breast cancer cell lines were analyzed by multiple factor analysis. (sentence)

variable and depends on the subtask. Figure 6 illustrates the difference between the two kinds of linking, and Table 9 shows examples from the F1000RD corpus.

4.5.2 Annotation: Explicit Links. The explicit linking layer in F1000RD was created in a two-step process. Based on an initial analysis of the corpus, we compiled a comprehensive list of targets-of-interest for explicit linking, including line numbers, page numbers,

Table 10
Distribution of explicit anchors and links in the F1000RD sample.

	lin	pag	col	par	quo	sec	fig	tab	box	ref	total
anchor	6	91	5	49	303	397	105	46	6	15	1,023
links	-	-	-	27	358	419	161	50	8	15	1,038

columns, paragraphs (e.g., “second paragraph”), quotes, sections, figures, tables, and references to other work. Two authors of the study manually annotated a random subset of 1,100 peer review sentences with explicit link anchors and their types, reaching 0.78 Krippendorff’s α on the general anchor/non-anchor distinction. Because of the good agreement, the rest of the data was manually annotated by one of the authors, who then also assigned target ITG nodes for each of the detected explicit anchors. Annotation was supported by a simple regular-expression-based parser, reaching 0.77 and 0.64 F1 score for explicit link anchor and target identification on our annotated data, respectively. The regular-expression-based approach failed in cases such as unspecific quotes, imprecisions in the review texts (e.g., spelling errors), and other edge cases not handled by the rigid, hand-coded system.

Table 10 shows the distribution of explicit anchor types and links in the resulting annotation layer: As we can see, explicit links are extensively used for referencing paper sections (sec), followed by literal quotes (quo) and figures (fig). Explicit links to lines (lin) and columns (col) are rare as the publication format in F1000Research is generally one-column and does not include line numbers. Page anchors (pag) are frequent—yet publications are only split into pages during PDF export; page numbers are not encoded in the source JATS XML files and thereby can’t be linked to.

4.5.3 Annotation: Implicit Links. Compared with explicit links, a major challenge in implicit link annotation is the absence of both an overt intertextuality marker on the anchor side, and a clear attachment point on the target side. This requires the annotator to simultaneously perform anchor identification and linking on two potentially large, previously unseen texts. In a range of pilot studies, we have attempted to separate the task into anchor identification and linking, similar to explicit link annotation; however, our experiments have demonstrated low agreement on anchor identification solely based on peer review text in G^R . We have thereby opted for a joint annotation scenario with linking formulated as sentence pair classification task: Given a pair of sentence nodes n_i^R from the peer review and n_j^P from the paper, the annotators needed to decide whether a link between the two sentences can be drawn.

While allowing the annotators simultaneous access to both texts, a pairwise classification setup inherently produces a large number of candidate pairs $|N^R| \times |N^P|$, most of which are irrelevant. To remedy this, similar to Radev, Otterbacher, and Zhang (2004) and Mussmann, Jia, and Liang (2020), we have implemented an annotation assistance system that given a review sentence n_i^R presents the annotator with a suggestion set S_i^P that consists of m most similar paper sentences $n_j^P \in N^P$ which are subsequently annotated as linked or non-linked to the review sentence. To diversify the suggestion set, we construct it by aggregating the rankings from multiple similarity functions: For our annotation study we have used the BM-25 score (Robertson and Zaragoza 2009) as well as cosine similarity between the review and paper sentence encoded using SentenceBERT (Reimers and Gurevych 2019) and Universal Sentence Encoder (Cer et al. 2018).

When there were overlaps in the highest-ranked sentences from the different methods, the highest-ranked sentence not yet in the list of selected sentences was chosen iteratively for each method.

Given the suggestions ($m = 5$), the annotators labeled the resulting sentence pairs according to the guidelines that specified the definition of a link and provided examples. Importantly, the annotators were asked to use paragraph level as the highest possible target granularity to avoid excessive linking of sentences that refer to the paper under review in general (e.g., “*The overall writing style is good*”). Given the guidelines and the annotation support system, after four hours of training two annotators labeled the F1000RD study sample. Even with annotation support, the task of implicit linking has proven substantially more challenging than pragmatic tagging, requiring twice as much time per review to produce annotations. The annotators have attained moderate (Landis and Koch 1977) agreement of 0.53 Krippendorff’s α , and the resulting 21,289 labeled sentence pairs for 4,819 review sentences make up the implicit linking layer of the F1000RD dataset. Unlike for pragmatic tagging, given the moderate agreement we decided not to perform adjudication and enforce a single gold standard annotation, and instead release the two sets of labels produced by the annotators separately, similar to Chandrasekaran et al. (2020) or Fornaciari et al. (2021).

4.5.4 Analysis. Our experiments demonstrate the large difference in complexity between the annotation of explicit and implicit links: Whereas explicit linking can be performed semi-automatically, implicit linking requires extensive annotation assistance and presents a major conceptual challenge. Determining the source of complexity in annotation of implicit links is hard: Related annotation studies by Chandrasekaran et al. (2020), Coffee et al. (2012), Wang et al. (2020) do not report inter-annotator agreement and don’t investigate the factors contributing to the disagreements. We have identified three such possible factors and conducted additional experiments to investigate their impact.

Domain Expertise. Our corpus includes scientific papers and their peer reviews from a wide range of domains, which might pose challenges both due to academic writing style and due to the domain knowledge required. Although the high agreement on the pragmatic tagging might suggest that our annotators are not affected by the domain, linking might require more intimate knowledge of the subject compared to pragmatic tagging. To investigate the impact of domain expertise on annotation performance, we recruited two additional expert annotators with strong medical background and conducted the implicit linking annotation study on a subset of F1000RD peer reviews and papers in the medical domain (case gateway), using the same protocol and guidelines as our main study. We then compared the agreement of the experts in their domain of expertise and in other domains: If the disagreements are indeed due to the lack of domain knowledge, we expect to see a higher agreement between the expert annotators in their domain of expertise. As Table 11 shows, this is not the case: Moreover, we observe *lower* agreement levels between the experts than between our main annotators who received additional training, annotated more data, and participated in the conceptualization of the task. This suggests that broad domain knowledge plays a secondary role in annotator agreement for the implicit linking task.

Subjectivity. Another potential source of low agreement is task subjectivity. Because implicit linking involves joint decision-making on anchor identification and target identification, it is vulnerable to disagreements on both whether the review sentence should

Table 11

Agreement statistics for implicit linking. The first column specifies the sets of annotations that are compared, where (A, B) are the annotations from the main study, (A_{re} , B_{re}) are from the re-test, and (C, D) from the experts. *o*, *m*, *n*: Agreement overall, on *medical*, on *non-medical*.

ann	#sent	#item	α
full			
A, B	4,819	21,289	0.53
re-test			
A, B			0.53
A_{re} , B_{re}	809	3,630	0.56
A, A_{re}			0.51
B, B_{re}			0.58
expert			
A, B	720	3,236	0.56^o / 0.59^m / 0.55^n
C, D			0.48^o / 0.48^m / 0.47^n

be linked at all, and what paper sentence it should be linked to. To study the effect of subjectivity, we conducted an additional experiment where, after a substantial period of time (2 months), the main annotators re-labeled a subset of documents from the F1000RD corpus. Such test-retest scenario allows us to measure not only the agreement levels between the two annotators, but also the self-agreement level. A moderate inter-annotator agreement with high self-agreement would signal that the task is conceptually easy, but highly subjective, since the annotators would disagree with each other, but conform to their own earlier decisions. As our results in Table 11 show, this is not the case: Whereas the inter-annotator agreement is slightly improved in the retest, the self-agreement values stay in the same moderate-agreement range as in the main annotation study. This suggests that the disagreements on implicit linking do not stem from the task subjectivity per se.

Task Conceptualization. The lack of domain expertise effect and the moderate self-agreement point to task definition as a potential target for further scrutiny. Intended as exploratory, our study deliberately left room for interpretation of the linking task. If high inter-annotator consistency is the goal, a stricter boundary between links and non-links would perhaps lead to higher agreement—and we see theoretical works in intertextuality theory as a promising source of inspiration for such delineation. We discuss other promising alternatives to our task definition and annotation procedure in Section 5.

4.6 Version Alignment

Change is a fundamental property of texts: While any text undergoes modifications as it is created, digital publishing makes it possible to amend texts even after publication, and this capability is widely used, from news to blog posts to collaborative encyclopediae. Academic publishing and text-based collaboration that surrounds peer review provide an excellent opportunity to study document change over time—yet, although some tools like Google Docs and Overleaf make it possible to trace the document changes on character level, in most cases only major revisions of texts are exchanged; whereas some publishers require the authors to describe the changes,

Downloaded from http://direct.mit.edu/col/article-pdf/48/4/949/2061780/col_a_00455.pdf by guest on 23 December 2023




Figure 7 Version alignment, only potential edges for a single node shown for clarity. The ILP formulation penalizes aligning nodes that belong to different node types (a) and uses node similarities (b) to construct a global alignment that maximizes total similarity (c) under constraints.

those descriptions are rarely enforced, not standardized, and not guaranteed to be complete. This makes it hard to verify whether the reviewers’ expert feedback has been addressed, and to find out which parts of the manuscript are new and require attention; the performance of the ad-hoc solutions like `diff-match-patch`¹² on manuscript version alignment in the academic domain has not been systematically assessed.

4.6.1 Task and Annotation. Motivated by this, we have conducted a study in automatic revision alignment of scientific manuscripts in the F1000RD sample. For simplicity, we cast the task as one-to-one ITG node alignment and only consider paragraph- and section-level alignment. Under those considerations, given two manuscript versions $G^{t+\Delta}$ and G^t , we aim to create a new set of edges $e(n_i^{t+\Delta}, n_j^t) \in E$ that signify the correspondences between revisions. Inspired by the work in graph-based annotation projection (Fürstenau and Lapata 2009), we formulate our approach as a constrained optimization via integer linear programming (ILP): Given the set of nodes $(n_1^{t+\Delta}, n_2^{t+\Delta}, \dots, n_i^{t+\Delta}) \in G^{t+\Delta}$ and $(n_1^t, n_2^t \dots n_j^t) \in G^t$, we define a binary variable $x_{i,j}$ that takes the value of 1 if $n_i^{t+\Delta}$ is aligned to n_j^t , that is, if we draw an edge $e(n_i^{t+\Delta}, n_j^t)$, and 0 otherwise. We then seek to maximize the objective $\sum_i \sum_j x_{i,j} * score(n_i^{t+\Delta}, n_j^t)$ under the one-to-one alignment constraints $\forall j \sum_i x_{i,j} \leq 1$ and $\forall i \sum_j x_{i,j} \leq 1$. As part of the scoring function, we use Levenshtein ratio and word overlap to compute similarity between ITG nodes sim ; in addition, we penalize the alignment of nodes that have different node types (e.g., paragraph and section title): $score(n_i^{t+\Delta}, n_j^t) = 0$ if $type(n_i^{t+\Delta}) \neq type(n_j^t)$, else $sim(n_i^{t+\Delta}, n_j^t)$. The result of the graph alignment is a set of cross-document edges connecting the nodes from $G^{t+\Delta}$ to G^t . Figure 7 illustrates our version alignment approach. The nodes in $G^{t+\Delta}$ with no outgoing alignment edges are considered added; the nodes in G^t with no incoming edges from the future version of the document, are considered deleted.

The alignments produced this way¹³ were evaluated by three expert annotators, who were presented with aligned manuscript revisions and asked to judge the correctness of the node pairings. Although it would be possible to include node splitting and

12 <https://github.com/google/diff-match-patch>.

13 We used the PuLP library <https://coin-or.github.io/pulp> as our ILP solver.

Table 12
Number of versions per paper in the full F1000RD corpus and the sample, by domain (Section 4.2).

#versions	full	sample	diso	iscb	rpkg	case	scip
1	3,743	50	10	18	14	–	8
2	1,353	105	22	9	12	42	20
3	243	17	5	4	2	3	3
4+	47	–	–	–	–	–	–

node merging into our objective by modifying the ILP constraints, this remained beyond the scope of our illustrative study.

4.6.2 Analysis. As the earlier Table 3 demonstrates, we oversample documents with revisions. A more detailed view in Table 12 shows that a significant number of documents in F1000RD undergo at least one revision before acceptance. We see that the number of revisions varies by the gateway, reflecting potential differences in publishing and reviewing practices across research communities.

For a more in-depth analysis, we focus on the differences between original submissions and their first revisions, as they are most numerous in our data. As Figure 8 demonstrates, in most cases the second revision of the manuscript contains more nodes, signifying incremental growth of the text in response to reviewer feedback. We note that the lack of change in the total number of nodes does not mean the lack of edits—those edits simply do not affect the document structure. These results reflect another property of the F1000Research publishing workflow: The lack of a formal page limit allows the authors to add information without the need to fit and restructure the publication, which is often not the case with other publishers. This observation highlights




Figure 8
Difference in the number of nodes between submission and first revision in the F1000RD sample, total number of papers. Negative values mean that the revised version is shorter than the original submission.

Table 13

Version alignment precision for submission and first revision in the F1000RD study sample, precision without exact matching, and the proportion of documents with perfect alignment. Only paragraph and section title-level nodes are considered.

	precision	precision w/o exact	perfect alignment
Levenshtein distance (norm)	0.982	0.966	0.713
Word overlap	0.985	0.973	0.746

the importance of taking the publishing workflow into account when working with the revision data.

Finally, we turn to the quality of our approach to automatic version alignment. Table 13 shows the performance of our proposed alignment method when validated by the three expert annotators. As we can see, our simple ILP-based aligner reaches good alignment precision independent of the similarity metric; at the same time we note the lower number of documents where *all* paragraph and title nodes have been correctly aligned, indicating room for improvement. While paragraph-level alignment is sufficient for our joint modeling study, more fine-grained analysis might be desirable for other tasks—we discuss this and other further directions for the study of version alignment in Section 5.

4.7 Joint Modeling

Together, pragmatic tagging, linking, and version alignment allow us to cover one full reviewing and revision cycle of a common peer reviewing process. Each of the analysis types allows us to answer practical questions relevant to text based collaboration: Once automated, pragmatic tagging can help to quickly locate relevant feedback and analyze the composition of the reviewing reports; linking enables navigation between peer reviews and their papers; and version alignment makes possible easy comparison of document revisions. The joint representation of the three phenomena by ITG allows us to explore additional questions that provide deeper insights into text-based collaboration during peer review.

Data Preparation. For each paper in the F1000RD sample, we construct an ITG for the paper itself, its reviews, and its first revision. We aggregate pragmatic tagging annotations for the review, explicit and implicit linking edges between the review and the paper, as well as version alignment edges. We limit implicit links to the cases where both annotators agree. Note that while pragmatic tagging and implicit linking are performed at the sentence level, version alignment happens at the paragraph and section level, and the granularity of explicit links might vary. The ability of the ITG to handle different granularities allows us to integrate these annotations: While pragmatic tagging remains on the sentence level, linking annotations are propagated to paragraph granularity to make them interoperable with version alignments (see Figure 2).

Why Reviewers Link. Reviewers use implicit and explicit links to refer to the paper they discuss—yet one might expect that the use of links is not uniform. We use our data to investigate the interaction of review pragmatics and linking, and find out why reviewers link. As Figure 9 shows, there is indeed a clear dependency between linking




Figure 9 Pragmatics of peer reviews and the linking behavior.

behavior and pragmatics of peer review reports: for example, Weaknesses get linked almost twice as often as Strengths, potentially reflecting the fact that whereas the praise tends to address the work in general, criticism is likely to point to particular locations in the manuscript that need improvement. Recap is rarely linked to the work explicitly, but often rephrases and summarizes the paper content, producing an implicit link. Todo sentences are only linked to the manuscript in one-third of cases, pointing to a potential target for improvement in the reviewing guidelines—yet, we note that a Todo might be related to a prior Weakness statement that contains a link, reflecting a limitation of our sentence-level approach to pragmatic tagging, which we revisit later.

What Gets Discussed. Just as the link anchors depend on peer review pragmatics, one could assume that link targets depend on the pragmatics of the papers under review—each section is not equally likely to be addressed by the reviewers. Unlike free-form peer review reports, the pragmatics of research publications is explicitly signalled by paper structure, and for this experiment we have mapped the section titles of the F1000RD sample to the few groups common to academic publishing—title, abstract, introduction, methods, results, discussion, and conclusions. We then calculated the number of incoming links that each of those sections accumulates, normalized by the number of times the section name appears in F1000RD. Figure 10 shows the results of our analysis along with the pragmatic category distribution on the peer review side. As we can observe, the most linked-to sections are Results and Methods—yet, the distribution of linked-from statements from the peer reviews also differs depending on the paper section and the type of linking—for example, while Abstracts are a frequent target of implicit Recap, Results are often explicitly referenced by Weakness sentences. Although deeper analysis of those interactions lies beyond the scope of this work, we note the depth of analysis enabled by combining different intertextual signals.

What Triggers Change. Finally, the version alignment annotations allow us to study the effects of peer reviewing on manuscript revision. We first analyze the interaction between linking and changes: As Figure 11 (left) shows, while most paper sentences

Downloaded from http://direct.mit.edu/col/article-pdf/48/4/949/2061780/col_a_00455.pdf by guest on 23 December 2023

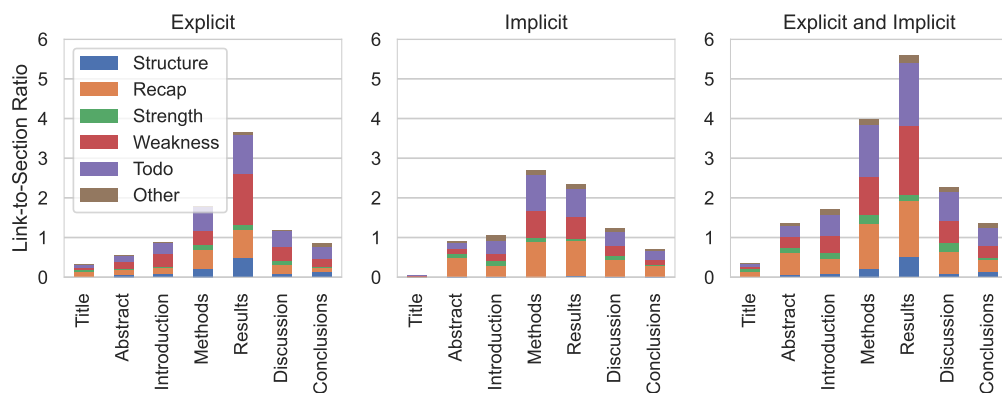


Figure 10 Pragmatic categories, links, and paper sections. The number of links per general paper section is normalized by the number of times the section appears in F1000RD.




Figure 11 Left: Linking and node change in F1000RD data. Right: Distribution of review sentence pragmatics for sentences linked to the paper paragraphs that were updated in a revision.

are not linked from the peer review, the linked ones (i.e., the ones discussed by the review) are almost twice as likely to be changed in a subsequent revision (probability of change is 0.55 and 0.30 for paragraphs with and without incoming links, respectively). On the review side, we analyze what kinds of peer reviewer statements tend to trigger change (Figure 11, right), and find that among the linked-to paragraphs of the paper, the most impact on the manuscript revision comes from the *Todo* sentences, followed by *Recap* and *Weakness*—whereas other pragmatic categories are only responsible for few revisions. We note that our analysis only provides a coarse picture of reviewing and revision behavior due to the differences in granularity and the intrinsic issues associated with sentence-level modeling of pragmatics; for example, a high number of changes triggered by the supposedly neutral *Recap* sentences points to potential limitations of our model. Although our proposed framework allows a more fine-grained investigation, here we leave it for the future.

5. Discussion

Having reviewed the instantiation of our proposed text-based collaboration model in the peer reviewing domain, we now take a step back and discuss the implications of our

results specific to peer review, and in general. Our proposed model is centered around three core intertextual phenomena: pragmatic tagging is an instance of architextuality—the relationship between a text and its abstract genre prototype; implicit linking is a reflection of metatextuality—the relationship between a text and its commentary; finally, version alignment taps into the paratextuality—by modeling the relationship between a text and its versions. Our proposed general view on intertextuality in NLP allowed us to systematize cross-document relationships along several basic axes (explicitness, type and granularity) surfacing connections between previously unrelated phenomena and highlighting the gaps that our subsequent study in the peer reviewing domain aims to fill. Modeling different phenomena within a joint framework makes it easy to study interactions between different types of intertextuality, and allows us to explore new research questions. To support our study, we propose a novel graph-based data model that is well suited for multidocument, relation-centered analysis on different granularity levels. The model generalizes over the common structural elements used to encode texts and can be easily extended to new domains and source types—while the current study has focused on the data from F1000Research, our current implementation includes converters for Wikipedia, with the support for more text sources on the way. Our studies in peer reviewing discourse provide valuable, open, cross-domain datasets for the development of novel NLP applications for peer reviewing support. We deliver insights that can help shape future studies in modeling text-based collaboration, and we now briefly outline our main takeaways.

Data Model. The main motivation behind our proposed graph-based data model is to facilitate modeling of intertextual relations on different granularity levels—for example, from sentence to paragraph. Yet, a general representation of document structure and other non-linguistic signals that surround texts has other potential uses, including language model fine-tuning and easy experimentation in cross-domain modeling; in addition, a graph-based data model allows preserving multimodal content—like tables or figures—by encapsulating it in the corresponding nodes, available for future processing. The further development of the ITG is to address three key challenges. As support for more input formats will be added, the data model definition is likely to be refined. To allow massive language model pre-training and fine-tuning, computational overhead associated with additional processing needs to be addressed. Finally, new conceptual solutions are necessary to efficiently utilize the additional information encoded by ITGs in modern Transformer-based NLP set-ups addressing particular end tasks.

Pragmatic Tagging. Our approach to pragmatic tagging as a classic sentence labeling task has shown good results—the annotations are reliable and the schema provides good coverage for the discourse analysis of peer reviews. Although our label set is tailored towards peer reviews, alternative schemata can be developed or adapted from prior work to cover new text genres, for example, scientific publications or news. Within the peer reviewing domain, our analysis suggests some additional pragmatic classes to be included in future versions of the labeling schema, most prominently, the performatives (“*I thereby accept this paper*”) and verbose confidence assessments (“*I’m not an expert in this area, but...*”). While in this work we resorted to sentence-level granularity for simplicity, a *clause* provides a perhaps more natural unit of analysis for discourse tagging—for example, as in Hua et al. (2019). We leave exploring the effects of granularity on annotation speed and quality for the future, noting that the ITG data model readily supports sub-sentence granularities.

Linking. Our analysis of linking behavior in the peer reviewing domain has revealed that the two distinct linking mechanisms—explicit and implicit linking—differ substantially in terms of annotation and processing complexity. Explicit linking can be largely tackled by simple rule-based approaches—facilitated by the ability of ITG to represent different structural elements of text. Yet, implicit linking requires more than that, as the annotation task has proven hard, reaching agreement levels similar to related studies (Radev, Otterbacher, and Zhang 2004) and other challenging tasks like argumentation mining (Habernal and Gurevych 2017). Our additional experiments involving re-test and expert annotations revealed that neither domain expertise nor subjectivity appear to affect the agreement—yet the lack of such comparisons in related work prevents us from making general conclusions about the nature of the task. We note that even our domain-expert annotators were neither authors nor reviewers of the annotated publications. Soliciting implicit linking annotations from the reviewers and authors presents a viable, yet organizationally challenging, alternative to producing labels in external annotation studies. From the annotation perspective, we deem it promising to explore the effects of the annotation interface and suggestion methods on the annotation quality and efficiency. From the task definition perspective, a promising alternative to our binary labeling schema is a decompositional approach where instead of producing a single binary label, annotators would answer a range of questions about the relationship between the anchor sentence and the target document, in the spirit of decompositional semantics (White et al. 2020) applied to challenging sub-sentential phenomena. We leave this exploration to future work.

Version Alignment. We have proposed a simple ITG alignment technique that does not require supervision, and—as our results demonstrate—provides good-quality paragraph-level alignments of F1000Research article revisions. Our proposed method is flexible and allows incorporating additional logical restrictions into the alignment process via ILP constraints. Yet it is important to note that the version alignment problem is far from solved: Despite the high precision scores, only 70% of the documents in our study were aligned perfectly. Moreover, the revision practices might vary substantially among research communities and publishing platforms. This might make the direct application of our proposed method problematic—for example, as F1000Research does not put a size limitation on the publications, many revisions grow incrementally—yet a page limit might potentially increase the number of modifications, as well as splitted and merged paragraphs, which are currently not supported by our aligner. Furthermore, while paragraph-level granularity has proven sufficient in our analysis, it might be insufficient for other applications. We deem it important to determine the parameters that affect revision practices across application scenarios and communities, and to collect diverse corpora of long-scope document revisions to support further investigation of the version alignment task.

Joint Modeling. Our final study in joint modeling of peer reviewing discourse has demonstrated the advantages of an integrated approach to text-based collaboration within the proposed data model. While the results reported here are illustrative and much deeper analysis is possible, we note that some limitations of the proposed approaches only become evident when the tasks are considered jointly. Our analysis of the reviewers' linking behavior revealed that an additional mechanism for modeling linking scope could be beneficial—although only mentioned in a single sentence, a link might in fact connect the whole subsequent segment of a text to a location in another text. Whether

linking scope should be modeled as part of pragmatic tagging and segmentation or as a separate information layer remains an open question. The optimal granularity level for the analysis of linking and revision behavior demands future investigation, as well.

Future Work. The core deliverables of this work are F1000RD—the first NLP corpus in open, journal-style, multidomain peer review—and the first instantiation of the proposed intertextual framework in the peer reviewing domain coupled with the annotation guidelines for the three novel intertextual NLP tasks. Our results create new opportunities for both basic and applied research. From the basic research perspective, apart from the many open methodological challenges outlined above, F1000RD enables the cross-domain study of peer reviewing behavior across different communities represented at F1000Research. It can be enriched with further annotations and extended to incorporate more gateways and more data types available at the F1000Research platform, including author responses and version amendment notes. Our general-purpose annotation protocols can be used to manually enrich peer reviewing data from other sources and conduct same kind of analysis on a sample from another reviewing environment: We highlight that although Section 2.1 discusses the sources of peer reviewing data used in NLP *so far*, there exist many alternative publishing platforms like PLOS ONE¹⁴ that open parts of the editorial process to the public, creating new opportunities for the computational study of text-based collaboration in peer review.

From an applied perspective, F1000RD can be readily used to develop NLP models to assist peer reviewing and editorial work, for example, by giving reviewers real-time feedback on the pragmatics of their reports, helping authors locate referenced passages in the paper, and helping editors ensure that a new paper revision addresses the concerns raised by the reviewers. A wide spectrum of questions to be tackled by ongoing and future work in this area include finding the optimal NLP approach to automate pragmatic tagging, linking, and version alignment; integrating the resulting models into real-world reviewing environments; measuring the effects of NLP-powered assistance on reviewing and editorial work; as well as extrapolating these findings to the text-based collaboration scenarios beyond academic peer review.

6. Conclusion

Text-based collaboration is at the core of many processes that shape the way we interact with information in the digital age. Yet the lack of general models of collaborative text work prevents the systematic study of text-based collaboration in NLP. To address this gap, in this article we proposed a model of text-based collaboration inspired by related work in intertextuality theory, and introduced three general tasks that cover one full cycle of document revision: pragmatic tagging, linking, and version alignment. To support our study, we developed the Intertextual Graph—a generic data model that takes the place of the ad-hoc plain text-based representation and is well suited for describing long documents and cross-document relations. We investigated the application of the proposed model to peer reviewing discourse and created F1000RD—the first clearly licensed, multidomain NLP corpus in open post-publication peer review. Our annotation studies revealed the strengths and weaknesses of our proposed approaches to pragmatic tagging, linking, and version alignment, and allowed us to determine promising directions for future research. While our proposed framework

¹⁴ <https://journals.plos.org/plosone>.

for analysis of text-based collaboration in NLP is joint—covering different aspects of this collaboration—a further necessary prerequisite for generalization is its unification across applications and domains. Thus, along with refining the task definitions and developing NLP models for performing the annotation tasks automatically, we deem it crucial to expand our proposed framework to new scenarios—by creating new data and incorporating existing data sources from other key domains of text-based collaboration like Wikipedia, news, online discussion platforms, and others.

Acknowledgments

This work has been funded by the German Research Foundation (DFG) as part of the PEER project (grant GU 798/28-1), the European Union under the Horizon Europe grant #101054961 (InterText), and the LOEWE Distinguished Chair “Ubiquitous Knowledge Processing” (LOEWE initiative, Hesse, Germany). We would like to thank Kateryna Shutiuk, Hussain Kamran, Jekaterina Kuznecova, Leonard Niestadtötter, and Mario Hambrecht for their help during the annotation studies and piloting for this work.

References

- Afrin, Tazin and Diane Litman. 2018. Annotation and classification of sentence-level revision improvement. In *Proceedings of the Thirteenth Workshop on Innovative Use of NLP for Building Educational Applications*, pages 240–246.
- Anjum, Omer, Hongyu Gong, Suma Bhat, Wen-Mei Hwu, and JinJun Xiong. 2019. PaRe: A paper-reviewer matching approach using a common topic space. 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 518–528. <https://doi.org/10.18653/v1/D19-1049>
- Beltagy, Iz, Matthew E. Peters, and Arman Cohan. 2020. Longformer: The long-document transformer. *arXiv:2004.05150*.
- Broich, Ulrich, editor. 1985. *Intertextualität: Formen, Funktionen, anglistische Fallstudien*. Number 35 in *Konzepte der Sprach- und Literaturwissenschaft*. Niemeyer, Tübingen. <https://doi.org/10.1515/9783111712420>
- Caciularu, Avi, Arman Cohan, Iz Beltagy, Matthew Peters, Arie Cattan, and Ido Dagan. 2021. CDLM: Cross-Document language modeling. In *Findings of the Association for Computational Linguistics: EMNLP 2021*, pages 2648–2662. <https://doi.org/10.18653/v1/2021.findings-emnlp.225>
- Cer, Daniel, Yinfei Yang, Sheng-yi Kong, Nan Hua, Nicole Limtiaco, Rhomni St. John, Noah Constant, Mario Guajardo-Cespedes, Steve Yuan, Chris Tar, Brian Strope, and Ray Kurzweil. 2018. Universal sentence encoder for English. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 169–174. <https://doi.org/10.18653/v1/D18-2029>
- Chandrasekaran, Muthu Kumar, Guy Feigenblat, Eduard Hovy, Abhilasha Ravichander, Michal Shmueli-Scheuer, and Anita de Waard. 2020. Overview and insights from the shared tasks at scholarly document processing 2020: CL-SciSumm, LaySumm and LongSumm. In *Proceedings of the First Workshop on Scholarly Document Processing*, pages 214–224. <https://doi.org/10.18653/v1/2020.sdp-1.24>
- Cheng, Liying, Lidong Bing, Qian Yu, Wei Lu, and Luo Si. 2020. APE: Argument pair extraction from peer review and rebuttal via multi-task learning. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 7000–7011. <https://doi.org/10.18653/v1/2020.emnlp-main.569>
- Coffee, Neil, Jean-Pierre Koenig, Shakthi Poornima, Christopher W. Forstall, Roelant Ossewaarde, and Sarah L. Jacobson. 2012. The Tesserae Project: Intertextual analysis of Latin poetry. *Literary and Linguistic Computing*, 28(2):221–228. <https://doi.org/10.1093/l1c/fqs033>
- Cohan, Arman, Waleed Ammar, Madeleine van Zuylen, and Field Cady. 2019a. Structural scaffolds for citation intent classification in scientific publications. 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 3586–3596. <https://doi.org/10.18653/v1/N19-1361>
- Cohan, Arman, Iz Beltagy, Daniel King, Bhavana Dalvi, and Dan Weld. 2019b.

- Pretrained language models for sequential sentence classification. 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 3693–3699.
- Cohan, Arman, Sergey Feldman, Iz Beltagy, Doug Downey, and Daniel Weld. 2020. SPECTER: Document-level representation learning using citation-informed transformers. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 2270–2282. <https://doi.org/10.18653/v1/2020.acl-main.207>
- Dasigi, Pradeep, Kyle Lo, Iz Beltagy, Arman Cohan, Noah A. Smith, and Matt Gardner. 2021. A dataset of information-seeking questions and answers anchored in research papers. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 4599–4610. <https://doi.org/10.18653/v1/2021.naacl-main.365>
- Daxenberger, Johannes and Iryna Gurevych. 2013. Automatically classifying edit categories in Wikipedia revisions. In *Proceedings of the 2013 Conference on Empirical Methods in Natural Language Processing*, pages 578–589.
- Daxenberger, Johannes and Iryna Gurevych. 2014. Automatically detecting corresponding edit-turn-pairs in Wikipedia. In *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 187–192. <https://doi.org/10.3115/v1/P14-2031>
- Devlin, Jacob, 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.
- Dycke, Nils, Edwin Simpson, Ilia Kuznetsov, and Iryna Gurevych. 2021. Ranking scientific papers using preference learning. *arXiv:2109.01190*.
- Esteva, Andre, Anuprit Kale, Romain Paulus, Kazuma Hashimoto, Wenpeng Yin, Dragomir Radev, and Richard Socher. 2021. COVID-19 information retrieval with deep-learning based semantic search, question answering, and abstractive summarization. *NPJ Digital Medicine*, 4(1):68. <https://doi.org/10.1038/s41746-021-00437-0>, PubMed: 33846532
- Fornaciari, Tommaso, Alexandra Uma, Silviu Paun, Barbara Plank, Dirk Hovy, and Massimo Poesio. 2021. Beyond black & white: Leveraging annotator disagreement via soft-label multi-task learning. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 2591–2597. <https://doi.org/10.18653/v1/2021.naacl-main.204>
- Forstall, Christopher W. and Walter J. Scheirer. 2019. *What is Quantitative Intertextuality?* Springer International Publishing, Cham. <https://doi.org/10.1007/978-3-030-23415-7.1>
- Fürstenau, Hagen and Mirella Lapata. 2009. Graph alignment for semi-supervised semantic role labeling. In *Proceedings of the 2009 Conference on Empirical Methods in Natural Language Processing*, pages 11–20.
- Genette, Gérard. 1997. *Palimpsests: Literature in the Second Degree*. Lincoln: University of Nebraska Press.
- Ghosal, Tirthankar, Sandeep Kumar, Prabhat Kumar Bharti, and Asif Ekbal. 2022. Peer review analyze: A novel benchmark resource for computational analysis of peer reviews. *PLOS ONE*, 17(1):1–29. <https://doi.org/10.1371/journal.pone.0259238>, PubMed: 35085252
- Ghosal, Tirthankar, Rajeev Verma, Asif Ekbal, and Pushpak Bhattacharyya. 2019. DeepSentiPeer: Harnessing sentiment in review texts to recommend peer review decisions. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 1120–1130. <https://doi.org/10.18653/v1/P19-1106>
- Habernal, Ivan and Iryna Gurevych. 2017. Argumentation mining in user-generated web discourse. *Computational Linguistics*, 43(1):125–179. <https://doi.org/10.1162/COLI.a.00276>
- Head, Andrew, Kyle Lo, Dongyeop Kang, Raymond Fok, Sam Skjonsberg, Daniel S. Weld, and Marti A. Hearst. 2021. Augmenting scientific papers with just-in-time, position-sensitive definitions of terms and symbols. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems, CHI '21*, pages 1–18.
- Hua, Xinyu, Mitko Nikolov, Nikhil Badugu, and Lu Wang. 2019. Argument mining for understanding peer reviews. 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 2131–2137. <https://doi.org/10.18653/v1/N19-1219>
- Jiang, Siyuan, Ameer Armaly, and Collin McMillan. 2017. Automatically generating commit messages from diffs using neural machine translation. *arXiv:1708.09492*. <https://doi.org/10.1109/ASE.2017.8115626>
- Kang, Dongyeop, Waleed Ammar, Bhavana Dalvi, Madeleine van Zuylen, Sebastian Kohlmeier, Eduard Hovy, and Roy Schwartz. 2018. A dataset of peer reviews (PeerRead): Collection, insights and NLP applications. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 1647–1661. <https://doi.org/10.18653/v1/N18-1149>
- Kardas, Marcin, Piotr Czapla, Pontus Stenetorp, Sebastian Ruder, Sebastian Riedel, Ross Taylor, and Robert Stojnic. 2020. AxCell: Automatic extraction of results from machine learning papers. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 8580–8594. <https://doi.org/10.18653/v1/2020.emnlp-main.692>
- Kristeva, Julia. 1980. Word, dialogue, and novel. In Leon S. Roudiez, editor, *Desire in Language: A Semiotic Approach to Literature and Art*. New York: Columbia University Press, pages 64–91.
- Kwiatkowski, Tom, Jennimaria Palomaki, Olivia Redfield, Michael Collins, Ankur Parikh, Chris Alberti, Danielle Epstein, Illia Polosukhin, Jacob Devlin, Kenton Lee, Kristina Toutanova, Llion Jones, Matthew Kelcey, Ming-Wei Chang, Andrew M. Dai, Jakob Uszkoreit, Quoc Le, and Slav Petrov. 2019. Natural questions: A benchmark for question answering research. *Transactions of the Association for Computational Linguistics*, 7:453–466. <https://doi.org/10.1162/tacl.a.00276>
- Landis, J. Richard and Gary G. Koch. 1977. The measurement of observer agreement for categorical data. *Biometrics*, 33(1):159–174. <https://doi.org/10.2307/2529310>, PubMed: 843571
- Lauscher, Anne, Goran Glavaš, and Simone Paolo Ponzetto. 2018. An argument-annotated corpus of scientific publications. In *Proceedings of the 5th Workshop on Argument Mining*, pages 40–46.
- Liu, Yinhan, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. RoBERTa: A robustly optimized BERT pretraining approach. *arXiv:1907.11692*, <abs/1907.11692>.
- Lo, Kyle, Lucy Lu Wang, Mark Neumann, Rodney Kinney, and Daniel Weld. 2020. S2ORC: The semantic scholar open research corpus. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 4969–4983. <https://doi.org/10.18653/v1/2020.acl-main.447>
- Luan, Yi, Luheng He, Mari Ostendorf, and Hannaneh Hajishirzi. 2018. Multi-task identification of entities, relations, and coreference for scientific knowledge graph construction. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 3219–3232.
- Mann, William C. and Sandra A. Thompson. 1988. Rhetorical structure theory: Toward a functional theory of text organization. *Text – Interdisciplinary Journal for the Study of Discourse*, 8(3):243–281. <https://doi.org/10.1515/text.1.1988.8.3.243>
- Marcus, Mitchell P., Mary Ann Marcinkiewicz, and Beatrice Santorini. 1993. Building a large annotated corpus of English: The Penn Treebank. *Computational Linguistics*, 19(2):313–330. <https://doi.org/10.21236/ADA273556>
- Maziero, Erick Galani, Maria Lucia del Rosario Castro Jorge, and Thiago Alexandre Salgueiro Pardo. 2010. Identifying multidocument relations. In *Proceedings of the 7th International Workshop on Natural Language Processing and Cognitive Science*, pages 60–69.
- Mimno, David and Andrew McCallum. 2007. Expertise modeling for matching papers with reviewers. In *Proceedings of the 13th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, KDD '07*, pages 500–509.
- Moreno, Laura, Gabriele Bavota, Massimiliano Di Penta, Rocco Oliveto, Andrian Marcus, and Gerardo Canfora. 2017. ARENA: An approach for the automated generation of release notes. *IEEE Transactions on Software Engineering*, 43(2):106–127. <https://doi.org/10.1109/TSE.2016.2591536>
- Mussmann, Stephen, Robin Jia, and Percy Liang. 2020. On the importance of adaptive data collection for extremely imbalanced

- pairwise tasks. In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 3400–3413. <https://doi.org/10.18653/v1/2020.findings-emnlp.305>
- Nivre, Joakim, Marie-Catherine de Marneffe, Filip Ginter, Yoav Goldberg, Jan Hajič, Christopher D. Manning, Ryan McDonald, Slav Petrov, Sampo Pyysalo, Natalia Silveira, Reut Tsarfaty, and Daniel Zeman. 2016. Universal Dependencies v1: A multilingual treebank collection. In *Proceedings of the Tenth International Conference on Language Resources and Evaluation (LREC'16)*, pages 1659–1666.
- Nury, Elisa and Elena Spadini. 2020. From giant despair to a new heaven: The early years of automatic collation. *IT - Information Technology*, 62(2):61–73. <https://doi.org/10.1515/itit-2019-0047>
- Nye, Benjamin, Ani Nenkova, Iain Marshall, and Byron C. Wallace. 2020. Trialstreamer: Mapping and browsing medical evidence in real-time. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics: System Demonstrations*, pages 63–69. <https://doi.org/10.18653/v1/2020.acl-demos.9>, PubMed: 34136886
- Radev, Dragomir. 2000. A common theory of information fusion from multiple text sources step one: Cross-document structure. In *1st SIGdial Workshop on Discourse and Dialogue*, pages 74–83. <https://doi.org/10.3115/1117736.1117745>
- Radev, Dragomir, Jahna Otterbacher, and Zhu Zhang. 2004. CST bank: A corpus for the study of cross-document structural relationships. In *Proceedings of the Fourth International Conference on Language Resources and Evaluation (LREC'04)*, pages 1783–1786.
- Reimers, Nils and Iryna Gurevych. 2019. Sentence-BERT: Sentence embeddings using Siamese BERT-networks. 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 3982–3992. <https://doi.org/10.18653/v1/D19-1410>
- Robertson, Stephen and Hugo Zaragoza. 2009. The probabilistic relevance framework: Bm25 and beyond. *Foundations and Trends in Information Retrieval*, 3(4):333–389. <https://doi.org/10.1561/15000000019>
- Stab, Christian and Iryna Gurevych. 2014. Identifying argumentative discourse structures in persuasive essays. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 46–56.
- Stab, Christian and Iryna Gurevych. 2017. Parsing argumentation structures in persuasive essays. *Computational Linguistics*, 43(3):619–659. <https://doi.org/10.1162/COLI.a.00295>
- Steyer, Kathrin. 2015. Irgendwie hängt alles mit allem zusammen – Grenzen und Möglichkeiten einer linguistischen Kategorie ‘Intertextualität’. Textbeziehungen. Linguistische und literaturwissenschaftliche Beiträge zur Intertextualität. Stauffenburg, Tübingen, pages 83–106.
- Teufel, Simone. 2006. Argumentative zoning for improved citation indexing. *Computing Attitude and Affect in Text: Theory and Applications*, pages 159–169. https://doi.org/10.1007/1-4020-4102-0_13
- Teufel, Simone, Advait Siddharthan, and Dan Tidhar. 2006. Automatic classification of citation function. In *Proceedings of the 2006 Conference on Empirical Methods in Natural Language Processing*, pages 103–110.
- Thelwall, Mike, Eleanor-Rose Papas, Zena Nyakoojo, Liz Allen, and Verena Weigert. 2020. Automatically detecting open academic review praise and criticism. *Online Information Review*, 44(5):1057–1076. <https://doi.org/10.1108/OIR-11-2019-0347>
- Thorne, James, Andreas Vlachos, Christos Christodoulopoulos, and Arpit Mittal. 2018. FEVER: A large-scale dataset for fact extraction and VERification. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 809–819. <https://doi.org/10.18653/v1/N18-1074>
- Wadden, David, Shanchuan Lin, Kyle Lo, Lucy Lu Wang, Madeleine van Zuylen, Arman Cohan, and Hannaneh Hajishirzi. 2020. Fact or fiction: Verifying scientific claims. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 7534–7550. <https://doi.org/10.18653/v1/2020.emnlp-main.609>
- Wang, Qingyun, Qi Zeng, Lifu Huang, Kevin Knight, Heng Ji, and Nazneen Fatema

- Rajani. 2020. ReviewRobot: Explainable paper review generation based on knowledge synthesis. In *Proceedings of the 13th International Conference on Natural Language Generation*, pages 384–397.
- White, Aaron Steven, Elias Stengel-Eskin, Siddharth Vashishtha, Venkata Subrahmanyam Govindarajan, Dee Ann Reisinger, Tim Vieira, Keisuke Sakaguchi, Sheng Zhang, Francis Ferraro, Rachel Rudinger, Kyle Rawlins, and Benjamin Van Durme. 2020. The Universal Decompositional Semantics dataset and Decomp toolkit. In *Proceedings of the 12th Language Resources and Evaluation Conference*, pages 5698–5707.
- Yang, Diyi, Aaron Halfaker, Robert E. Kraut, and Eduard H. Hovy. 2017. Identifying semantic edit intentions from revisions in Wikipedia. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing, EMNLP 2017*, pages 2000–2010.
- Yuan, Weizhe, Pengfei Liu, and Graham Neubig. 2021. Can we automate scientific reviewing? *arXiv:2102.00176*.
- Zhang, Fan and Diane Litman. 2015. Annotation and classification of argumentative writing revisions. In *Proceedings of the Tenth Workshop on Innovative Use of NLP for Building Educational Applications*, pages 133–143. <https://doi.org/10.3115/v1/W15-0616>