Breaking Ties: Some Methods for Refactoring RST Convergences

Andrew Potter, Cameron Lindsey
University of North Alabama
Florence, Alabama, USA
apotter1@una.edu, clindsey4@una.edu

Abstract

Among the set of schemata specified by Rhetorical Structure Theory is a pattern known variously as the request schema, satellite tie, multisatellite nucleus, or convergence. The essential feature of this schema is that it permits multiple satellites to attach to a single nucleus. Although the schema has long been considered fundamental to RST, it has never been subjected to detailed evaluation. This paper provides such an assessment. Close examination shows that it results in structures that are ambiguous, disjoint, incomplete, and sometimes incoherent. Fortunately, however, further examination shows it to be unnecessary. This paper describes the difficulties with convergences and presents methods for refactoring them as explicit specifications of text structure. The study shows that convergences can be more clearly rendered not as flat relational conjunctions, but rather as organized expressions of cumulative rhetorical moves, wherein each move asserts an identifiable structural integrity and the expressions conform to specifiable scoping rules.

1 Introduction

Among the fundamentals of Rhetorical Structure Theory is a set of abstract relational patterns, one of which permits linking multiple satellites to a single nucleus. This pattern has been referred to variously as the *request schema*, *satellite tie*, *multisatellite nucleus*, and *convergence*. Opinions vary as to the validity of this schema. It was baselined in Mann and Thompson's (1988) original RST specification, endorsed by Carlson and Marcu (2001) and by Stede, Taboada, and Das (2017), but rejected by Zeldes (2017, 2023) as well as by Shahmohammadi and Stede (2024). However, none of these studies, either pro or con, has provided any rationale for the position taken. That leading researchers in RST disagree over a fundamental property of satellite-nucleus structures and yet provide no grounded account for the discourse pattern is concerning. This has implications for RST parsers and other applications, as well as for the theory's efficacy in describing discourse. Moreover, if the schema is to be rejected, what criteria are to be applied in determining the structures used in its stead? Thus, the motivation for this study is not merely theoretical, but is of practical interest as well.

Our primary finding is that the convergence is problematic. Convergences are found to be functionally ambiguous, implicitly disjoint, structurally incomplete, and sometimes incoherent. These problems arise, due not to any issue in the texts under analysis, but as a result of the analytical abstraction itself. The good news, however, is that convergences are also entirely unnecessary. Moreover, their avoidance results in greater functional specificity than would otherwise be attainable. Since it is possible to produce high-resolution analyses without using convergence, and since the status of convergence is questionable at best, methods for refactoring RST convergent structures should be of interest.

The process for inferring asymmetric structures from convergences is referred to as *refactoring*. We implemented and assessed three refactoring algorithms. The first of these we call the *chaining* method. Using chaining, the convergence is restructured as a sequence of relations in which each successive satellite feeds into the satellite next closest to the nucleus. Although such patterns do occur in RST discourse, refactoring convergences as chains did not usually result in plausible RST analyses. The second algorithm is limited to convergences in which all satellites share the same relation, so that they can be restructured as a *list*. The third algorithm nests satellites according to precedence so that relations take on the appearance of the rungs of a *ladder*. This method is the most flexible of the three and produces high-resolution structures that are capable of handling complex convergences with satellites straddling the nucleus. This approach is the principal focus of this paper. The study included the refactoring of 279 RST analyses, each containing at least one convergent structure. The analyses were drawn from the literature as well as from several online corpora.

Refactored analyses can be used in the same ways as other RST analyses. The difference is that refactored analyses are more precise in their structural definitions, thus providing higher fidelity input for downstream processes, such as parsing and summarization. And, to the extent that RST diagrams are in and of themselves useful, refactored analyses provide more informative visualizations.

2 Related Work

Two closely related theories of text organization are foundational to this research. The first of these is, obviously, *Rhetorical Structure Theory*. RST is a conceptual framework that explains clausal text organization in terms of the way the text spans comprising the text relate to one another (Mann & Thompson, 1988). RST postulates a small number of patterns (or schemas) for defining the structural possibilities among spans, and it defines a set of rhetorical relations for use when applying a schema to the text spans. Second, *Relational Propositions*, also developed by Mann and Thompson (1983), are implicit coherence-producing assertions that serve to bind together explicit parts of a text and are essential to the effective functioning of the text. RST analyses and relational propositions are isomorphic. For every relation in the rhetorical structure, a corresponding relational proposition is asserted (Mann & Thompson, 1986, p. 268). Potter (2019a, 2023) extended Man and Thompson's theory of relational propositions to make it interchangeable with RST. This included development of a predicate notation conformant with the Python programming language. This enables the treatment of RST analyses as data, code, or diagrams, serving as an enabling research technology. The refactoring methods used here were implemented as Python scripts, all working within this framework. A key addition to this is a program developed to reverse the process, transforming relational propositions back into RST, storing the result as RS3/XML files, the format developed by O'Donnell (1997) and the *de facto* data sharing format for RST. An overview of end-to-end refactoring is shown in Figure 1. This software was used to generate the refactored analyses directly from the originals.

The use of convergent structures has been widely accepted among many RST researchers. They appear regularly in the seminal publications of Mann and Thompson, and they have been used without question in numerous other research publications (e.g., Abelen, Redeker, & Thompson, 1993; Fiacco, Jiang, Adamson, & Rosé, 2022; Potter, 2019b; Wang, Wu, & Cui, 2020). A few researchers have, however, sown seeds of doubt. Egg and Redeker (2008) suggested that, to the extent that discourses could be analyzed as trees, the possibility of convergent structures seemed unlikely. Their assumption was that relations interpreted as tree nodes could have no more than one parent. They further suggested that convergent structures might be genre-dependent, perhaps found mainly in fund-raising letters, as found in studies by Mann et al. (1992) and Abelen et al. (1993). However, the corpora examined in this study suggest that the use of convergences is in fact broadly applied across multiple genres. About a third of the analyses in the Potsdam Corpus (Stede & Neumann, 2014), and one fourth of the analyses in the Online Learning Corpus (Potter, 2008) contain at least one convergence. Since this structural pattern is at best unnecessary (as this paper demonstrates), perhaps its presence or absence has more to do with analyst preference than with any particular genre.

Carlson, Marcu, and Okurowski (2003) viewed RST convergent (as well as other) relations as static and discrete, such that each relation is to be understood in isolation from its neighbors. Thus they viewed satellites of convergences as independent or separate modifications to the nucleus. This would suggest that a convergence is merely a diagrammatic overlay of unrelated structures, rather than an integrated explanation of text organization. The interpretation described in this paper provides an alternative view, that the assertion of a relational proposition is a discursive event with an identifiable effect, such that multiple interrelated relational propositions will have a developing effect, and these effects are observable by means of refactoring. That this should be the case becomes clear when considering the incoherence that arises when inconsistent satellites are associated with a shared nucleus.

As far as we know, Zeldes (2023) and Shahmohammadi and Stede (2024) are the only researchers who have rejected convergences (or *satellite ties*, in Zeldes' terminology). Although Zeldes used convergences in earlier research (2016), more recent versions of the GUM corpus contain no convergences, but rather follow a strict hierarchical regimen (Zeldes et al., 2024).

3 The Corpus

As mentioned above, the corpus consists of 279 RST analyses drawn from a variety of research publications and open-source corpora. All analyses include at least one convergent relation. Some basic examples come from Mann and Thompson's early RST publications (1987, 1988). Others were chosen selectively, including papers by Mann, Thompson and Matthiessen (1992) and by Abelen, Redeker, and Thompson (1993). These include Satellite-Nucleus-Satellite patterns they identified as of special interest (discussed in Section 5.3). Matthiessen and Thompson's (1987) example of convergence as LIST motivated the evaluation given in Section 5.2. Others are from Carlson and Marcu (2001), Stede et al. (2017), Bateman (2001), 2008), two texts from the *Multilingual RST Treebank* (Iruskieta, da Cunha, & Taboada, 2014), and 219 analyses from Potter's (2008) *Online Learning Discussion* corpus.

4 The Refactoring Software

The software developed for this research leverages two open-source applications. The first of these converts RS3 formatted RST files into Pythonic relational propositions. The second evaluates nested relational propositions as Python expressions, with each relation defined as a function. Both algorithms are described in detail by Potter (2023, 2024a). We also developed a script that converts Pythonic relational propositions to RS3 format for end-to-end interchangeability. The software, as well as the corpus used in this study are as open source from GitHub. [1]

5 Problems with Convergence

There has been a longstanding tendency to treat RST analyses as top-down hierarchies. This follows from the orientation of the diagrams and has been reinforced by Marcu's strong nuclearity assumption. And yet, unless we wish to view satellites as dispensable, we must recognize that a nucleus obtains its stature from the satellites that support it. To understand what is *happening* in an RST diagram, then, it is appropriate to start at the bottom of the structure, with the outermost extremities, and *follow the arrows* through the series of relations to the ultimate locus of intended effect. Within the structure, span nodes function as precedence operators, nesting the relational propositions, one within one another, building outward (Potter, 2024b). This is what gives RST structures their tree-like appearance. In contrast to this, convergences have a flattening effect on the structure. A convergence asserts that multiple relational propositions terminate on a shared

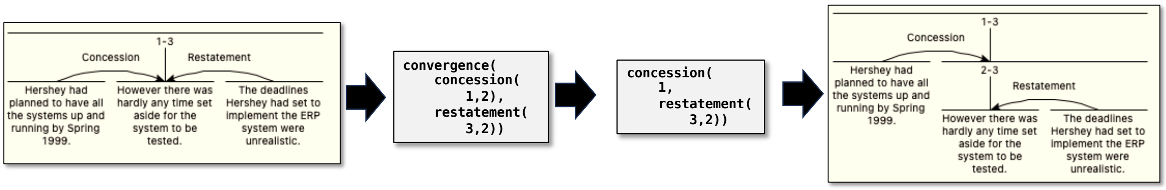

Figure 1. Automated Refactoring Process. (Adapted from Ducasse & Brown, 2023)

Ducasse and Brown (2023), Egg and Redeker (2008), and Zeldes (2017), whose evolving *GUM News Worship* analysis provides external validation of the ladder method described in Section 5.3. The critical mass of analyses come from open-source corpora. These include 12 analyses (translated into English) from the *Potsdam Commentary Corpus of German newspaper editorials* (Stede & Neumann, 2014), 14 messages from the *STS-Corpus* of listserv emails exchanged in a scholarly debate (Potter,

[1] All data and code are downloadable from https://github.com/anpotter/RST-Refactor

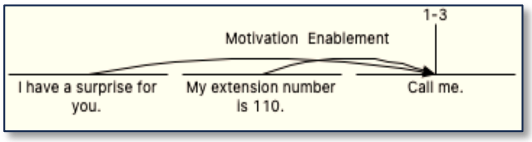

Figure 2. Convergence as a set of Distinct Rhetorical Moves (Mann & Thompson, 1986)

nucleus, with each satellite occupying a shared level of precedence. But this misrepresents what happens in discourse. Consider the convergence shown in Figure 2. It contains two rhetorical moves, one corresponding to relational proposition *motivation(1,3)* and the other to relational proposition *enablement(2,3)*. The intended effect of the first relational proposition is to *motivate* the reader to make the call, and the intended effect of the second is to *enable* the reader to make the call. These effects enact separate moves. That the writer has deemed both moves as necessary to achieve the shared effect informs us that their synergy is realized cumulatively, not simultaneously. This means that either *enablement* is subordinate to *motivation*, or *motivation* is subordinate to *enablement*. The precedence within the text is unspecified in the diagram.

Similar difficulties arise when satellites converge to the right of the nucleus. And if the satellites are repositioned so that one precedes the nucleus and another follows, although the disjuncture is avoidable, the ambiguity persists. Either the left-hand side takes precedence over the right, or the right over the left. In their discussion of what they called the *satellite-nucleus-satellite* (SNS) pattern Mann et al. (1992) proposed that convergences similar to this are argumentatively strategic. However, the strategy they describe, while perhaps evident in the text, is unsupported by the symmetry of the convergent structure. For RST to do its work, any such strategy must be reflected in the structure. But convergent structures *conceal* rather than *present* this structural functionality.

And this is only the beginning of the problem. Convergences such as MOTIVATION-ENABLEMENT may seem readily intuitive, with each proposition contributing to a shared intentionality. And convergences consisting of pilings-on of multiple instances of identical relations, such as repetitive ELABORATION relations likewise have intuitive appeal. Mann and Thompson's (1988) *Syncom* analysis is a good example of this. But this intuitiveness masks a difficulty. RST schemas place no constraints on what relations may be combined, allowing a mix of any sort whatsoever. This might at first seem reasonable, as it would seem the text under analysis should be the deciding factor. But in practice this results in some odd bedfellows.

Consider this example from Ducasse and Brown (2023), shown in Figure 3. As far as classic RST is concerned, there is no problem here. That there is a CONCESSION relation from unit 1 to 2 seems clear, as is the RESTATEMENT from 3 to 2. However, the analysis as a whole is problematic. As analyzed, the scope of the CONCESSION relation applies only to unit 2, and the scope of the RESTATEMENT also applies only to unit 2. But if unit 3 is indeed a RESTATEMENT of 2, the scope of the CONCESSION will be applicable to both 2 and 3. There is a structural discrepancy with the analysis arising, not from the choice of relations, but from the use of convergence. Disorganizations of this sort become more pronounced with complex structures. But RST has everything it needs to make the development explicit. Convergence is not among them. And that is a problem that can be addressed by refactoring.

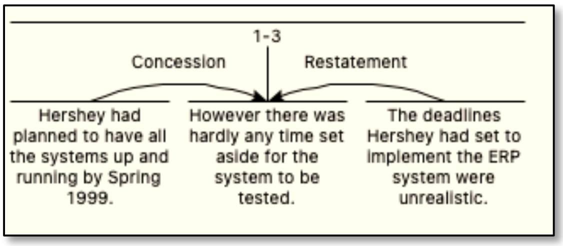

Figure 3. Convergence as a Problematic Combination of Moves (RST excerpt from Ducasse and Brown (2023))

6 Refactoring Convergences

Refactoring consists of a set of methods for restructuring convergences to conform to satellite-nucleus, nucleus-satellite, or multinuclear patterns. As shown in Figure 4, *chaining* links successive satellites, each attaching to its successor in a stair-step pattern; *listing* groups convergent satellites as multinuclear LIST; and *laddering* links satellites to a cumulative nuclear span using the most direct path available. A scoping heuristic is provided for determining satellite subordination when the convergence straddles its nucleus. All three methods conform to the minimal RST constraints for completeness, connectedness, uniqueness, and adjacency. Beyond that, their differences are considerable.

6.1 Chaining

For *chaining* to render a plausible analysis it would be necessary for the convergent relations to form a succession of relational dependencies, one providing functional support for the next. While this sometimes occurs in discourse, in the convergences examined in this study, this seemed to be at best only weakly supported, with lists and ladders consistently resulting in more plausible interpretations.

6.2 Listing

Matthiessen and Thompson (1987) proposed that when the satellites of a structure “function as *co-equal* realizations of a single relation” the LIST relation could be used for what would otherwise be specified as a convergence (p. 288). A similar stance was adopted by Shahmohammadi and Stede (2024). While it is unspecified as to what the specific criteria for *co-equality* are, the RST definition of the LIST relation requires that the items must be at least somehow *comparable* to one

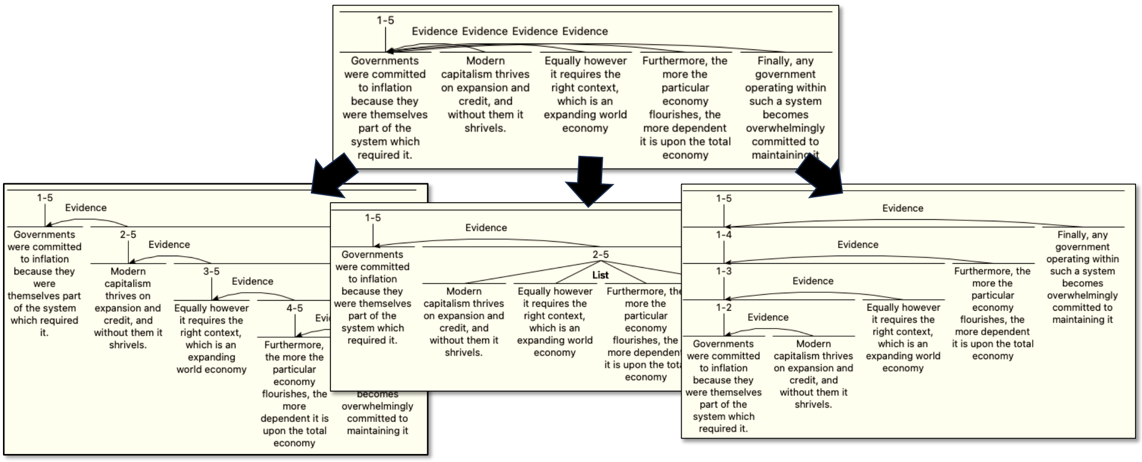

Figure 4. Chaining, Listing, and Laddering (Refactored from Martin, 1992)

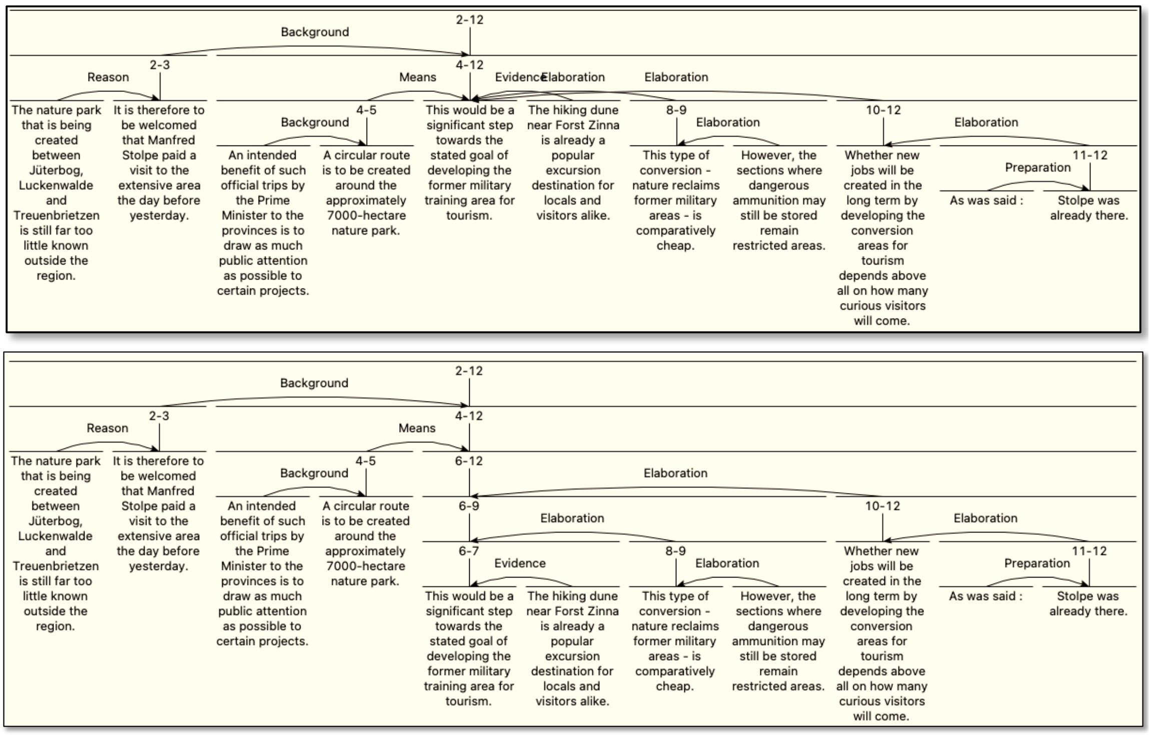

Figure 5. Laddering a Complex Convergence (Refactored from Stede & Neumann, 2014e)

another. In any case, LIST can be applied to single-relation convergences only, and all members must be on the same side of the nucleus. Further, what might conform relationally as a LIST relation may go beyond comparable membership, such that the succession of members exhibits intentional development or argumentative accrual. As shown in Figure 4, each successive satellite builds on the effect of its predecessors, as Martin (1983) and Bateman (2001) observed. While listing tends to produce concise, readable diagrams, this can result in a loss of structural information. This information can be readily rendered using the ladder method.

6.3 Laddering

Laddering refactors a convergence by extending its nuclear span node to connect each relation successively, such that precedence decreases inversely to satellite distance from the nucleus. When the convergence occurs to the left of the nucleus, the successive relations *close in* on the nucleus. When the convergence occurs on the right-hand side, where the relations *build out* from the nucleus. This enables expression of the rhetorical structure as a fully articulated relational proposition. Laddering readily generates plausible solutions for convergences when all satellites attached on the same side of the nucleus. The closer a satellite is to its nucleus, the higher its precedence. However, the situation becomes more interesting when the convergence includes satellites on both sides of the nucleus. Figure 5 shows a functionally overloaded nucleus. The nucleus of the convergence is simultaneously a situation that might be realized (MEANS), asserted to be a belief the reader will hopefully accept (EVIDENCE) and a situation in need of additional information (ELABORATION). These relations all reach the nucleus without interaction among one another. Granted, the text may be such that this assortment of relations comes into play. But in a convergence, they are depicted as simultaneous. That they may be diagrammed as such, however, does not make it so. As each relation effects the nucleus, the nucleus is functionally modified. What is missing is the structural path through which this happens. Clearly there is an ongoing development. And indeed, through refactoring, also shown in Figure 5, this process can be defined. The situation to be realized by the MEANS is assured and elaborated by the EVIDENCE and ELABORATION, but it is not the case that these assurances and

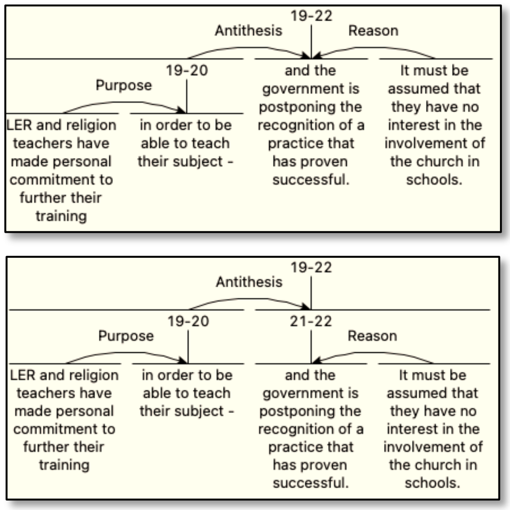

Figure 6. Using Scope Escalation to Determine Precedence (Refactored from Stede & Neumann, 2014e)

elaborations are applied to the MEANS. They constitute the nucleus of the MEANS. Indeed, by the time the MEANS reaches the unary nucleus, the EVIDENCE and ELABORATION will each have long since done their work. Convergences simply do not occur. The rhetorical development playing out among the MEANS, EVIDENCE, and ELABORATION relations is inaccessible in the convergent interpretation. It is at best a stand-in for a series of discursive moves.

When convergences straddle the nucleus, determining precedence for the order of moves becomes more interesting. The ordering depends on the combination of relations in use and requires *scope escalation*. To say that one relation has scope over another is equivalent to saying the other takes precedence over the one. Although scoping decisions for convergences that are restricted to one side of the nucleus are determined by their order of appearance, this determinant is unavailable when convergences straddle the nucleus. Either the left-hand side will take scope over the right, or the right over the left. In the convergence shown in Figure 6, either the ANTITHESIS will take scope over the REASON relation, or the REASON relation will take scope over the ANTITHESIS. The scope escalation procedure is used to determine which side will have precedence. In this example, ANTITHESIS has scope over REASON.

Inevitably, one move has scope over the other. The question is, how are scope escalations between straddling convergences to be determined? Mann et al. (1992) provide a clue. In their study of the rhetorical structure of a fund-raising letter, they briefly discussed the aforementioned *Satellite-Nucleus-Satellite* (SNS) pattern, which they considered to be unusual. An SNS consists of a nucleus flanked by two satellites of the same relation, as shown in Figure 7. Mann et al. (1992) proposed that this pattern implements a presentational mode of *leading up to the point, stating the point, and driving the point home*. Several instances of the pattern were found in their analysis. Abelen et al. (1993) also mention finding the pattern in their study of fund-raising letters, and they agreed with Mann et al.'s characterization of it as a presentational mode. That this interpretation is said to derive from the SNS pattern is a matter of interest. Looking at the structure, what we see are two identically related satellites straddling a nucleus. The higher order interpretation, that the left-hand side of the relation *leads up to a point* and that the right-hand side *drives it home* suggests a more complex structure, one that enacts the dynamic of the argument. While the convergent structure is symmetric, the interpretation of the structure is clearly asymmetric. There is more going on here than simple convergence. This can be leveraged for making scoping determinations.

A feature of the pattern is the repetition of the relation on the left and right sides. This repetition builds on the relational intention, fortifying the functionality. This being the case, a more descriptive fit would be for the left-hand side of the structure, leading up to the point, to be subordinated under the right-hand side, where the point is then driven home. This becomes clear when cycling through a corresponding relational proposition, combining the reenactment and compression algorithms defined by Potter (2024a). Figure 8 shows the progression of the refactored structure as it the steps through the moves identified by Mann et al. (1992). The process begins with the elementary propositions on the left-hand side, establishes the causal linkage, and then drives home the point using the right-hand side.

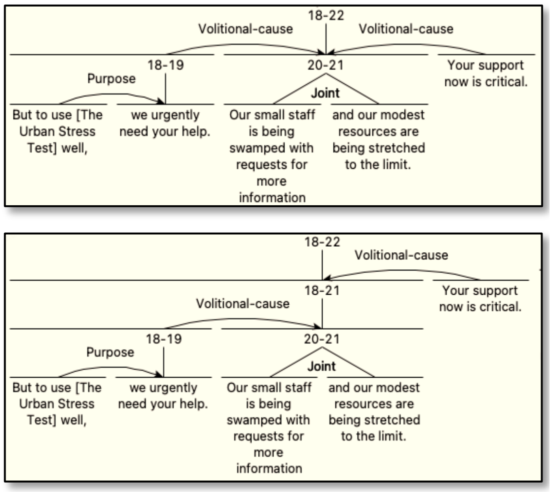

Figure 7. SNS Pattern and Refactored SNS Pattern from (Refactored from Mann, Matthiessen, & Thompson, 1992)

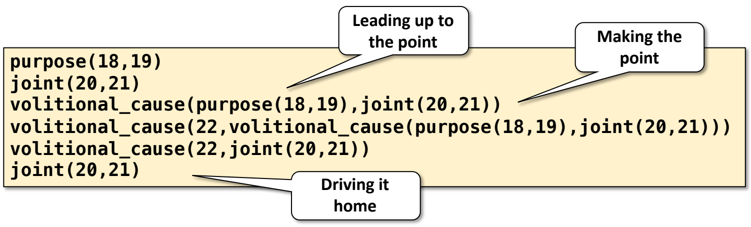

Figure 8. Cycling Through a Refactored SNS Structure

Mann et al. (1992) and Abelen et al. (1993) limited their identification of the SNS pattern to pairs of identical relations. This is more restrictive than necessary, both in terms of the number of satellites and with respect to identicality. When the number of satellites exceeds two, precedence takes care of itself, building out on the right-hand side, and closing in on the left-hand side. The limitation to identical relations is also unnecessary, because the effect need not be restricted to specific relations, but rather to relational intentionality. Relations of similar intentionality may be matched as readily as identical pairs. Thus, for refactoring purposes, the SNS pattern is extensible. With this in mind, a set of categories of relational effects were adapted from Stuart-Smith (2007), who developed as typology of rhetorical relations based on Systemic Functional Linguistics. The adapted categories of relational effects is shown in Table 1. When convergences belonging to the same category straddle the nucleus, the scope of the right-hand side is escalated over the left side. An exception to this is the SUMMARY relation. As a satellite restatement of its nucleus, its scope is escalated irrespective of the category of the left-hand side. Any convergence not meeting this criterion defaults to scope given to the left-hand side. Thirty-one examples of scoping were found in the corpus. This includes not only fund-raising letters, but also various news articles and online discussion items, suggesting the phenomenon may be generalizable beyond the fund-raising genre.

7 Conclusion

Precedence and scoping are fundamental to the representation of discourse processes. The order of moves depends on satellite positioning with respect to the nucleus and to neighboring satellites. The closer a satellite is to its nucleus, the higher its precedence. With satellites to the right of the nucleus, precedence decreases with each successive satellite. The contribution of each is evaluated in light of its predecessor. Satellites to the left of the nucleus gain precedence the closer they are to the nucleus. Like the right-hand side, the effects are cumulative, but rather than reinforcing, the effect is anticipatory, such that they lead up to and prepare for the locus of intended effect. The interchangeability between RST diagrams and relational propositions supports scalability, consistency, and reproducibility.

One might well ask, however, if convergences are so undesirable, how is it that RST has gotten along with them so well for so long? Even assuming refactoring is an improvement, is it significant? These questions go to the heart of what makes RST interesting. RST has been used in a wide range of applications, and yet these applications tend to make little use of the diagrams per se. Yet it is the diagrams that make RST distinctive. In this light, we can say that while an RST *analysis* is an articulation of the intentional structure of a discourse, an RST *diagram* is a point-by-point account of a discourse process. The process initiates with the outermost elementary relational propositions. These propositions join to form more complex expressions which ultimately specify the comprehensive discourse process terminating with the ultimate locus of intended effect. Thus, the process of reading of a diagram is, rather than a top-down activity, a bottom-up process, a process of following the arrows.

As we consider RST diagrams as pictures of processes, we raise the possibility of concurrent threads. Recent work by Zeldes et al. (2024) indicates that concurrent relations are common. Although little attention has been given to simultaneous analyses, might not the ambiguity of convergent structures suggest that simultaneity may be more common than has been hitherto supposed? Although the scoping criteria for determining precedence between left- and right-hand satellites appear generalizable, there might be sufficient exceptions to make simultaneity commonplace. Further studies in which RST diagrammatic semantics are a focus, not merely representations of analyses, would be useful here.

Category	**Relation**
Acceptance	EVIDENCE, JUSTIFY, REASON
Performance	ENABLEMENT, MOTIVATION
Comprehension	ELABORATION, SUMMARY, RESTATEMENT, EVALUATION, INTERPRETATION
Resistance	ANTITHESIS, CONCESSION
Causality	CIRCUMSTANCE, MEANS, NON-VOLITIONAL CAUSE, NON-VOLITIONAL RESULT, OTHERWISE, UNLESS, PURPOSE, SOLUTIONHOOD, UNCONDITIONAL, UNLESS, VOLITIONAL CAUSE, VOLITIONAL RESULT

Table 1. Categories of Relational Effects

Limitations

The most important (and obvious) delimitations in this study are those inherent in Rhetorical Structure Theory itself. While the primary claim of this research has been to question (and reject) a fundamental assumption of that theory, other assumptions remain unaddressed, such as the constraints of completedness, connectedness, uniqueness and adjacency. Other theories of discourse relations, to whatever extent the problem of convergence may or may not be relevant, have also been ignored. Another limitation of the paper concerns the size of the corpus. The topic requires that samples be limited to convergence-containing analyses. Although 279 analyses may seem small, note that this is only slightly smaller than the corpus used for the original development of Rhetorical Structure Theory (Mann, 2001).

References

Abelen, E., Redeker, G., & Thompson, S. (1993). The rhetorical structure of US-American and Dutch fund-raising letters. *Text - Interdisciplinary Journal for the Study of Discourse, 3*, 323-350.

Bateman, J. A. (2001). Between the leaves of rhetorical structure: Static and dynamic aspects of discourse organisation. *Verbum, 23*(1), 31-58.

Carlson, L., & Marcu, D. (2001, September). Discourse tagging reference manual. Retrieved from ftp://ftp.isi.edu/isi-pubs/tr-545.pdf

Carlson, L., Marcu, D., & Okurowski, M. E. (2003). Building a discourse-tagged corpus in the framework of rhetorical structure theory. In J. v. Kuppevelt & R. Smith (Eds.), *Current directions in discourse and dialogue*. Berlin: Springer.

Ducasse, A. M., & Brown, A. (2023). Rhetorical relations in university students' presentations. *Journal of English for Academic Purposes, 63*, 101251.

Egg, M., & Redeker, G. (2008). Underspecified discourse respresentation. In A. Benz & P. Kühnlein (Eds.), *Constraints in Discourse* (pp. 117-138). Amsterdam: John Benjamins.

Fiacco, J., Jiang, S., Adamson, D., & Rosé, C. (2022). Toward Automatic Discourse Parsing of Student Writing Motivated by Neural Interpretation. In *Proceedings of the 17th Workshop on Innovative Use of NLP for Building Educational Applications (BEA 2022)*. Seattle, Washington: Association for Computational Linguistics.

Iruskieta, M., da Cunha, I., & Taboada, M. (2014). A qualitative comparison method for rhetorical structures: identifying different discourse structures in multilingual corpora. *Language Resources and Evaluation, 49*, 263-309.

Mann, W. C. (2001, April 2). Authority to speak: The Justify relation -- some issues. *RSTlist.* Retrieved from http://listserv.linguistlist.org/pipermail/rstlist/2001-April/000091.html

Mann, W. C., Matthiessen, C. M. I. M., & Thompson, S. A. (1992). Rhetorical structure theory and text analysis. In W. C. Mann & S. A. Thompson (Eds.), *Discourse description: Diverse linguistic analyses of a fund-raising text* (pp. 39-78). Amsterdam: John Benjamins.

Mann, W. C., & Thompson, S. A. (1983). *Relational propositions in discourse*. Marina del Rey, CA: Information Sciences Institute.

Mann, W. C., & Thompson, S. A. (1986). Assertions from discourse structure. In *HLT '86: Proceedings of the workshop on strategic computing natural language* (pp. 257-270). Morristown, NJ: Association for Computational Linguistics.

Mann, W. C., & Thompson, S. A. (1987). *Rhetorical structure theory: A theory of text organization* (ISI/RS-87-190). Retrieved from Marina del Rey, CA:

Mann, W. C., & Thompson, S. A. (1988). Rhetorical structure theory: Toward a functional theory of text organization. *Text - Interdisciplinary Journal for the Study of Discourse, 8*(3), 243-281.

Martin, J. R. (1983). Conjunction: The Logic of English text. In J. S. Petöfi & E. Sozer (Eds.), *Micro and Macro Connexity of Texts* (pp. 1-72). Hamburg: Helmut Buske Verlag.

Martin, J. R. (1992). *English text: System and structure*. Philadelphia: John Benjamins.

Matthiessen, C. M. I. M., & Thompson, S. A. (1987). The structure of discourse and 'subordination'. In J. Haiman & S. A. Thompson (Eds.), *Clause combining in grammar and discourse* (pp. 275-329). Amsterdam: John Benjamins.

O'Donnell, M. (1997). RST-Tool: An RST analysis tool. In *Proceedings of the 6th European Workshop on Natural Language Generation*. Duisburg, Germany: Gerhard-Mercator University.

Potter, A. (2008). Interactional coherence in asynchronous learning networks: A rhetorical approach. *The Internet and Higher Education, 11*, 87–97.

Potter, A. (2019a). Reasoning between the lines: A logic of relational propositions. *Dialogue and Discourse, 9*(2), 80-110.

Potter, A. (2019b). The rhetorical structure of attribution. In A. Zeldes, D. Das, E. M. Galani, J. D. Antonio, & M. Iruskieta (Eds.), *Proceedings of the Workshop on Discourse Relation Parsing and Treebanking (DISRPT2019)* (pp. 38-49). Minneapolis, MN: Association for Computational Linguistics.

Potter, A. (2023). An algorithm for Pythonizing rhetorical structures. In S. Carvalho, A. F. Khan, A. O. Anić, Blerina Spahiu, J. Gracia, J. P. McCrae, D. Gromann, Barbara Heinisch, & A. Salgado (Eds.), *Language, data and* *knowledge 2023 (LDK 2023): Proceedings of the 4th Conference on Language, Data and Knowledge* (pp. 493-503). Vienna, Austria: NOVA CLUNL.

Potter, A. (2024a). An Algorithmic approach to analyzing rhetorical structures. In M. Strube, C. Braud, C. Hardmeier, J. J. Li, S. Loaiciga, A. Zeldes, & C. Li (Eds.), *Proceedings of the 5th Workshop on Computational Approaches to Discourse (CODI 2024)* (pp. 1-11). St. Julians, Malta: Association for Computational Linguistics.

Potter, A. (2024b). *Hiding in plain sight: Span nodes as first-class objects in RST*. Paper presented at the Beyond Words: Theoretical, Experimental, and Computational Approaches to Language, Contexts, and Modalities, Gothenburg, Sweden.

Shahmohammadi, S., & Stede, M. (2024). Discourse parsing for German with new RST corpora. In *Proceedings of the 20th Conference on Natural Language Processing (KONVENS 2024)* (pp. 65-74). Vienna, Austria: Association for Computational Linguistics.

Stede, M., & Neumann, A. (2014). Potsdam Commentary Corpus 2.0: Annotation for discourse research. In *Proceedings of the International Conference on Language Resources and Evaluation (LREC)* (pp. 925-929). Reykjavik: European Language Resources Association.

Stede, M., Taboada, M., & Das, D. (2017). *Annotation guidelines for rhetorical structure*. Retrieved from Potsdam and Burnaby:

Stuart-Smith, V. (2007). The Hierarchical Organization of Text as Conceptualized by Rhetorical Structure Theory: A Systemic Functional Perspective. *Australian Journal of Linguistics, 27*(1), 41-61.

Wang, Y., Wu, H., & Cui, G. (2020). Rhetorical structure analysis of prepared speeches and argumentative essays by Chinese advanced English learners. *Text & Talk, 40*(2), 219-240. doi:doi:10.1515/text-2020-2054

Zeldes, A. (2016). rstWeb – A browser-based annotation interface for Rhetorical Structure Theory and discourse relations. In *Proceedings of NAACL-HLT 2016 (Demonstrations)* (pp. 1-5). San Diego, California: Association for Computational Linguistics.

Zeldes, A. (2017). The GUM corpus: Creating multilayer resources in the classroom. *Language Resources and Evaluation, 51*(3), 581-561.

Zeldes, A. (2023, November 20). Rhetorical Structure Theory annotation and eRST.

Zeldes, A., Aoyama, T., Liu, Y. J., Peng, S., Das, D., & Gessler, L. (2024). eRST: A signaled graph theory of discourse relations and organization. *Computational Linguistics, 51*(1), 23-72.