



Who Wrote This Email? Topical Structure as a Predictor of Native, Non-Native, and AI-Generated Writing

Sahar Zahed Alavi 

Department of Foreign Languages, University of Bojnord, Bojnord, Iran

Received: 2026/07/22

Accepted: 2026/08/25

Abstract: The emergence of artificial intelligence (AI) has raised important questions about how AI-generated texts differ from human writing at the discourse level. While previous research has primarily focused on lexical and syntactic features, little attention has been paid to topical progression as a discourse marker of organization. Following the framework of systemic functional linguistics (Halliday, 2009), this study investigated whether native, non-native, and AI-generated English business emails differ systematically in their topical progression patterns and whether these topical progressions can predict text origin. Four topical progression patterns (i.e., Parallel, Sequential, Extended Parallel, and Extended Sequential) were analyzed using Poisson generalized linear models controlling for text length, followed by multinomial logistic regression for text classification. The results revealed selective rather than uniform differences across the three text types. Native texts demonstrated significantly greater use of Extended Parallel topical progression, whereas AI-generated texts exhibited higher frequencies of Extended Sequential progression. Non-native texts differed from native texts primarily through a lower frequency of Parallel topical progression, while Sequential progression showed no significant differences across the groups. The multinomial model achieved an overall classification accuracy of 85.6%, indicating that topical progression frequencies provided strong discriminatory power. Thus, it could be an interpretable linguistic framework for distinguishing AI-generated, native, and non-native writing.

Keywords: Topical Structure Analysis, Business Emails, Generative AI, Authorship Prediction.

* Corresponding Author.

Authors' Email Address:

Sahar Zahed Alavi (s.zahedalavi@ub.ac.ir)



Introduction

In today's international business and professional contexts, digital communication has become the dominant mode of interaction; business emails serve as one of the most widely used channels of immediate and asynchronous communication (Liu & Liu, 2023). Efficient business emails play an important role in clearly conveying information, facilitating decision-making, and promoting intercultural collaboration (Sujinpram & Wannaruk, 2026).

At the same time, the writing landscape has been transformed by the rapid adoption of generative AI tools, which can produce fluent drafts at speed (Elkhatat, 2023; Haleem, 2023). This shift has intensified concerns about how writing produced by different sources can be distinguished: when an email sounds polished, is it written by a native speaker, a non-native professional using English as an additional language, or an AI system? In response, many institutions have turned toward AI text detectors, but evidence suggests that these tools are not equally reliable for all writers (Gao et al., 2023). In particular, GPT detectors have been shown to frequently misclassify non-native English writing as AI-generated, raising serious fairness concerns and calling into question the wisdom of relying on detector scores in evaluative contexts (Liang et al., 2023). Hence, there is a need to identify characteristics that may differentiate native, non-native, and AI-generated writing without relying exclusively on automated detector scores.

Previous research has mainly approached this problem through lexical, stylistic, and syntactic characteristics. Studies comparing AI-generated and human writing, for example, have examined vocabulary use, sentence structure, syntactic complexity, and other surface-level stylistic characteristics (Herbold et al., 2023; Liao et al., 2023; Reviriego et al., 2024). Although these approaches have demonstrated that writing sources can exhibit different linguistic profiles, they focus primarily on individual linguistic forms or sentence-level structure and provide less information about how ideas and topics are organized across sentences (Tsai, 2025; Wang, 2024). This distinction is particularly relevant to business emails, where effective communication depends not only on appropriate vocabulary and grammatical form but also on the coherent development of information throughout the message.

Topical Structure Analysis (TSA) offers a complementary discourse-level perspective by examining how sentence topics develop across a text (He et al., 2009; Lautamatti, 1987; Saeed et al., 2022; Quinto, 2015). Rather than measuring which words writers use or how individual sentences are syntactically constructed, TSA captures patterns of topical continuity and change, thereby providing an observable measure of how information is organized across

sentences. It specifically examines whether writers tend to maintain a topic, move from one topic to another through newly introduced information, or return to a previously established topic. Because these patterns reflect the distribution and development of topical information across a text, they provide a potential means of investigating differences in discourse organization that may not be captured by vocabulary or sentence-level syntax alone (Saeed et al., 2022).

Despite increasing research on differences between AI-generated and human writing, discourse-level organization remains comparatively underexplored, particularly in business email contexts. Existing comparative studies have provided evidence from lexical and syntactic features (Herbold et al., 2023; Liao et al., 2023; Reviriego et al., 2024), while research on discourse and coherence has emphasized the importance of information organization in written communication (Incelli, 2013; Tsai, 2025; Wang, 2024). However, what remains unclear is whether topical progression (i.e., the observable development of sentence topics across a text) varies systematically among native-speaker, non-native-speaker, and AI-generated business emails. More specifically, it is not yet sufficiently established whether the frequencies of different topical progression patterns differ among these three writing-source groups or whether those patterns provide sufficient information to predict group membership.

Addressing this gap is theoretically and practically significant. Theoretically, examining topical progression across the three groups can extend research on writing-source variation from lexical and sentence-level characteristics to the organization of discourse, thereby providing evidence about how coherence and information flow are realized in different types of writing. Methodologically, it can establish whether TSA provides discriminating information beyond approaches that focus primarily on lexical or syntactic features. Practically, identifying discourse-level differences may contribute to a more nuanced understanding of automated and human judgments of writing source, particularly given concerns that conventional AI-detection systems may incorrectly classify non-native English writing as AI-generated (Liang et al., 2023). Therefore, the present study examines the topical structure of business emails written by native speakers, non-native speakers, and AI and tests both whether the groups presented in this study's corpus differ in their use of topical progression patterns and whether these patterns can predict group membership. The following two research questions are raised:

1. Are there statistically significant differences among native, non-native, and AI-generated business emails in terms of the frequencies of the four topical progression patterns after controlling for text length?
2. Can topical progression patterns predict whether a business email is written by a native speaker, a non-native speaker, or AI?

Theoretical Framework

Systemic Functional Linguistics (Halliday, 2009) and Topical Structure Analysis (Lautamatti, 1987) are the theoretical frameworks of the present study. Halliday pointed to the metafunctions of language, among which the textual metafunction is responsible for making the text coherent. In this system, Halliday (2009) distinguished theme (i.e., the point of departure of the message) from rheme (i.e., the remainder of the information developed from the theme). In declarative clauses, the theme typically occupies the initial position, establishing the perspective from which subsequent information is interpreted. This distinction provides a framework for describing how information is organized within clauses and how the point of departure of a message relates to the information that follows.

Lautamatti (1987) introduced Topical Structure Analysis (TSA), one of the earliest systematic methods for examining topic development across entire texts. Rather than focusing solely on clause-initial themes, TSA tracks the progression of sentence topics throughout text to show how writers maintain, extend, or shift topical focus. Lautamatti (ibid) argued that coherence depends on the logical continuity of topical development; it could be investigated through observable patterns of topic progression rather than subjective judgments of textual quality. Her work represented an important methodological advancement because it operationalized discourse coherence into identifiable analytical categories that could be applied consistently in different texts (Saeed et al., 2022).

The original TSA framework distinguishes three patterns of topical development. Parallel progression occurs when consecutive clauses maintain the same topic, reinforcing continuity and emphasizing the presented topic. Sequential progression arises when information introduced in the rheme of one clause becomes the topic of the following clause. Extended Parallel progression refers to the reintroduction of an earlier topic after one or more intervening clauses. Simpson (2000) added Extended Sequential progression, in which the information introduced in the previous rheme becomes the basis for the subsequent topic. The examples of the types of topical progression are presented below.

- (1) **We** have now received the remaining payment for the Robofil 310 machine. **We** will be packing it into the boxes this week. (Parallel progression, Email 28 written by a native speaker)
- (2) **The Texas desk** needs to create a new product that is similar to the Exxon Katy Phys product. **The product** will be identical to the Exxon Katy product except for the location. (Sequential progression, Email 9 written by a native speaker)
- (3) **We** will transfer the payment for you subsequently. Please start the process of manufacturing the machines. **We** also inform you that we have paid you Euro 50,000.00 on November 20 for our previous order. (Extended Parallel progression, Email 19 written by a non-native speaker)
- (4) We are pleased to inform you that we have prepared an information package about our new trading solution. Please let us know if you would like to arrange a meeting to discuss your business needs. **The information package** includes product specifications, competitive pricing, and implementation guidelines. (Extended Sequential progression, Email 9 generated by AI)

Literature Review

Business Emails as a Professional Genre

The expansion of international commerce has established business email as a key genre of English as a Business Lingua Franca (EBLF), through which interlocutors from diverse linguistic and cultural backgrounds communicate. Compared with traditional business letters, EBLF emails tend to be more personalized, conversational, and similar to spoken interaction (Roshid et al., 2018). However, this apparent informality does not mean that business emails lack systematic organization. Rather, writers must balance transactional purposes, interpersonal relationships, and culturally situated expectations within a concise written format.

In ESP research, genre analysis has become the dominant framework for investigating this organization. Building on Swales' (1990) concept of genre as a class of communicative events sharing common communicative purposes, researchers have primarily examined business emails in terms of rhetorical moves and steps. These studies indicate that business emails exhibit relatively stable communicative functions but considerable variation in how those functions are realized. For example, Mehrpour and Mehrzad (2013) found that Iranian and native English-speaking writers used the same four broad moves (i.e., establishing the negotiation chain, providing information, requesting action, and ending the message) but differed in their realization. Iranian writers drew explicit attention to messages, apologized,

and relied on collective pronouns more frequently, whereas native English-speaking writers adopted a more individualized style. Thus, shared move structures do not necessarily imply shared discourse practices; writers may fulfill the same communicative purposes through culturally different interpersonal and organizational choices.

This interaction between transactional and interpersonal purposes is also evident across different business-email subgenres. [Van Herck et al. \(2022\)](#) identified six moves in responses to customer complaints: opening, acknowledging the complaint, brand positioning, transactional complaint handling, concluding remarks, and closing. Their analysis challenges accounts that treat business correspondence as primarily transactional by showing that greetings, apologies, and expressions of empathy are integrated with explanations and problem resolution. Similarly, [Park et al. \(2021\)](#) found that American professionals generally used more supportive moves, indirect requests, and relationship-building strategies, whereas Korean professionals placed greater emphasis on directness and repeated requests. Taken together, these findings show that rhetorical organization is shaped not only by communicative purpose but also by politeness norms, face management, writer-reader relationships, and cultural expectations.

Nevertheless, the move-analytic focus of this literature presents an important limitation. Move analysis identifies the broad communicative functions performed in a message, but it does not necessarily explain how writers connect information from one sentence or clause to the next. Two emails may contain the same moves while differing in the continuity, repetition, or development of their topics. Consequently, the existing genre literature provides a less developed account of the local discourse progression. Examining topical structure can therefore complement move analysis by revealing how information is maintained and developed within and across rhetorical moves.

AI-Generated Writing and Professional Communication

The emergence of AI has transformed professional communication. AI tools such as ChatGPT, Gemini, and Microsoft Copilot are increasingly used in the workplace to draft emails, prepare reports, summarize documents, and support routine communication tasks ([AlAfnan, 2025](#); [Cardon et al., 2023](#); [Getchell et al., 2022](#); [Olszewski et al., 2026](#)). Their appeal lies in their potential to accelerate writing processes, reduce costs, and produce fluent, apparently high-quality content ([Haleem et al., 2023](#); [Papakonstantinidis, 2024](#)). Fluency, however, can obscure important differences between surface-level quality and discourse-level effectiveness. AI-generated texts may resemble human writing to the extent that even

experienced language professionals struggle to distinguish them from human writing. [Casal and Kessler \(2023\)](#) reported that expert reviewers correctly identified AI-generated research abstracts only 33.7% of the time. Interestingly, reviewers based their judgments on specificity, writing quality, and authorial voice, yet these criteria were unreliable indicators of AI authorship. This finding suggests that intuitive judgments based on overall quality or style may be insufficient and that more systematic linguistic indicators are needed.

Comparative research generally portrays AI-generated writing as structurally polished but also potentially standardized. [Herbold et al. \(2023\)](#), for example, compared argumentative essays written by students with those generated by ChatGPT-3 and ChatGPT-4. Expert evaluators rated AI-generated essays significantly higher across several dimensions of writing quality, including logical organization, vocabulary sophistication, and language complexity. Computational analyses likewise indicated greater syntactic complexity and structural consistency in the AI-generated essays. At the same time, the essays repeatedly used similar introductions, conclusions, and organizational templates across topics. AI's consistency may therefore function both as an advantage and as a possible marker of authorship: it can enhance perceived organization while limiting rhetorical flexibility.

Evidence concerning AI's lexical performance is less uniform and appears to depend on the model, task, and professional domain. Although AI-mediated texts may employ terminology consistently ([AlAfnan, 2025](#)), [Reviriego et al. \(2024\)](#) found that ChatGPT-3.5 produced fewer distinct lexical items than human writers, indicating lower lexical diversity despite high fluency. ChatGPT-4, however, approached or occasionally exceeded human performance, depending on the task. Accordingly, low lexical diversity cannot be treated as a stable or model-independent characteristic of AI writing. In medical writing, [Liao et al. \(2023\)](#) similarly found that AI-generated texts relied on more general vocabulary and fewer specialized lexical items than texts written by medical professionals. Human-authored texts contained more varied terminology, numerical information, and domain-specific expressions. Considered together, these studies distinguish lexical fluency from informational precision: AI may produce accessible and cohesive text without consistently reproducing the specificity associated with expert professional knowledge. They also caution against identifying AI authorship through isolated lexical features, because such features can change across model generations and communicative contexts.

The limitations of lexical and syntactic indicators have directed increasing attention toward discourse organization. [Yang et al. \(2024\)](#) found that ChatGPT-generated argumentative essays relied predominantly on constant thematic progression, in which the

same theme was maintained across successive clauses. Although this pattern supported grammatical cohesion, it also produced repetitive and list-like discourse. Human writers used linear progression more frequently, allowing information introduced in one clause to become the point of departure for the next and thereby creating more cumulative development. This contrast suggests that the distinction between human and AI writing may lie less in whether a text is cohesive than in how that cohesion is achieved: AI-generated texts may favor thematic stability, whereas human writers may permit more dynamic shifts in information flow.

Research on professional email provides related evidence. [Wilson and Rose \(2025\)](#) found that ChatGPT, Gemini, and human writers all reproduced the obligatory moves of business refusal emails. However, the AI systems overused particular rhetorical steps, including repeated expressions of appreciation, empathy, and politeness, whereas human writers adapted move sequencing more flexibly to interpersonal relationships and communicative contexts. These results parallel the findings from argumentative writing: across genres, AI appears capable of reproducing expected structures but may realize them with greater repetition and standardization. Nevertheless, the evidence remains limited because genre moves and thematic progression operate at different levels of discourse. Successful reproduction of obligatory moves does not establish that AI organizes topics within those moves in the same way as human writers.

The preceding literature reveals four related gaps. First, studies of business email have concentrated on rhetorical moves ([Mehrpour & Mehrzad, 2013](#); [Park et al., 2021](#); [Van Herck et al., 2022](#)), politeness strategies ([Qu, 2024](#)), and linguistic features ([Roshid et al., 2018](#)). Although these approaches explain what communicative functions emails perform and how interpersonal meanings are expressed, they provide limited insight into how topics are maintained or developed across clauses. Second, human-AI comparisons have focused mainly on argumentative essays ([Herbold et al., 2023](#); [Yang et al., 2024](#)), medical writing ([Liao et al., 2023](#)), and academic texts ([Reviriego et al., 2024](#)), leaving business emails comparatively underexplored. Third, thematic progression research suggests that human and AI texts differ in information flow, but thematic progression is not identical to TSA. It therefore remains necessary to determine whether TSA captures comparable authorship differences in professional communication. Fourth, existing studies have largely described differences between groups; they have seldom tested whether discourse-level features can statistically predict authorship, particularly across native English-speaking, non-native English-speaking, and AI-generated texts.

The present study addresses these gaps by applying TSA to business emails written by native English speakers, non-native Iranian speakers, and generative AI. Rather than treating topical patterns only as descriptive characteristics, the study examines their potential as predictors of authorship. Specifically, it compares the frequencies of four topical-progression patterns while controlling for text length and tests whether these discourse-level features can distinguish among the three authorship groups.

Methods

Design

This study used a quantitative corpus-based discourse analysis to examine the topical organization of business emails written by native speakers, non-native speakers, and AI.

Corpus

The corpus for the present study consisted of 180 business emails: 60 written by native speakers, 60 by non-native Iranian speakers, and 60 by AI. The native-speaker and non-native-speaker emails were selected from a business email corpus collected under the coordination of Shiraz University. The native/non-native classification in the source corpus was based on the linguistic and contextual information available to the corpus compilers in establishing the two human-written groups. The native-speaker group comprised business correspondence classified in the source corpus as having been produced by native English-speaking writers in English-speaking professional contexts in England and the United States, whereas the non-native-speaker group comprised correspondence classified in the source corpus as having been produced by non-native English-speaking Iranian writers working in Iranian professional contexts. This corpus has previously been used in discourse-analytic research. In particular, [Mehrpur and Mehrzad \(2013\)](#) used this corpus to analyze the moves in business emails.

For the present study, emails were selected based on their primary communicative function. Specifically, emails were included if their main purpose was to provide information or to request information, an action, or a favor. This operational definition follows the idea that communicative purpose is the defining characteristic of a genre ([Swales, 1990](#)). To ensure confidentiality, all addresses appearing in the To and From fields were deleted.

The third category of the corpus consisted of AI-generated business emails. To help AI generate comparative emails, the communicative contexts of 30 randomly selected emails written by native speakers and those of 30 randomly selected emails written by non-native speakers were given to ChatGPT-5 one by one with a prompt asking for writing a business

email on the proposed communicative context (see Appendix B). Each communicative context was entered in a separate interaction, and the same generation procedure and prompt were used for all 60 AI-generated emails. No additional instructions concerning topical progression, native/non-native writing characteristics, or the linguistic features under investigation were provided. The AI-generated emails were not edited, paraphrased, or grammatically corrected by the researchers after generation. The outputs were analyzed in the form in which they were generated.

Table 1 presents the features of the corpus, including the total number of words in emails, the total number of t-units in emails, and the average number of words per t-unit. As is evident, the emails written by native speakers are longer than those written by AI and non-native speakers. However, the number of t-units generated by AI is more than that of native and non-native speakers.

Table 1. Features of the Corpus

Units	Native speakers	Non-native speakers	AI
Total number of words	5502	2781	3441
Total number of t-units	414	301	450
Average number of words per t-unit	13.28	9.23	7.64

The data were analyzed using Topical Structure Analysis, proposed by [Lautamatti \(1987\)](#) to examine how topics were developed throughout the business emails. The analysis was carried out in several stages. First, the number of words and the number of t-units (i.e., any independent clause and all its required modifiers, or any non-independent clause punctuated as a sentence, or any imperative ([Schneider & Connor, 1990](#)) in each business email were determined. Then, the topic in each t-unit was identified. The identification of topics rested on the use of interpretation to determine what a t-unit was about ([Witte, 1983](#); [Schneider & Connor, 1990](#)) and the use of “as for” device at the beginning of the t-unit ([Huddleston, 1984](#)). Beginning a t-unit with an “as for” phrase and finding an appropriate object for it separated the t-unit’s topic from other possibilities. Then, different types of topic progression (Parallel, Sequential, and Extended parallel) in each email were explored through some guidelines excerpted from [Schneider and Connor \(1990, p. 427](#); see Appendix A). [Simpson \(2000\)](#) added the fourth type of topic progression, Extended Sequential Progression, in which a unit in the rheme of an earlier t-unit became the topic of a later t-unit after an intervening sequence.

Finally, when the progressions were plotted, the number of occurrences of each type of topic progression in each email was counted and tabulated. It should be noted that the three

groups were not identical in their total number of words. This difference resulted from the use of authentic business emails, which naturally varied in length across the source groups. Because the purpose of the study was to compare topical progression patterns across individual business emails rather than to compare aggregate word totals, the analysis was designed to account for differences in text length. In particular, the statistical analyses used the number of T-units as the exposure measure.

Data analyses were conducted using R software. To answer the first research question, separate Poisson generalized linear models (GLMs) were fitted for each topical structure type (Parallel, Sequential, Extended Parallel, and Extended Sequential). The dependent variable was the frequency count of each topical structure, and text type (native, AI-generated, and non-native) was entered as the predictor. Native texts served as the reference category. Because the texts differed in length, the logarithm of the number of t-units was included as an offset in each model. Consequently, the models estimated the rate of occurrence of each topical structure per t-unit rather than raw frequency counts, thereby controlling for variation in text length. To account for differences in the text lengths, the logarithm of the number of T-units in each email was included as an offset in each model. This specification estimates the expected rate of occurrence of each topical structure per T-unit rather than comparing raw frequencies, thereby controlling for variation in email length at the level at which topical progression was measured. In addition, since four separate models were estimated, Holm-adjusted p-values were applied to control the Type I error rate across multiple comparisons.

To answer the second research question, to examine whether topical structure frequencies predicted text type, a multinomial logistic regression analysis was conducted using R software. The dependent variable was text type, consisting of three categories (native, AI-generated, and non-native), with native designated as the reference category. The predictor variables were the frequencies per t-unit of the four topical structure types: Parallel, Sequential, Extended Parallel, and Extended Sequential. Frequencies were converted to rates by dividing the frequency of each topical structure by the total number of t-units in each text to control for differences in text length. Model performance was evaluated using a likelihood ratio test comparing the fitted model with the intercept-only model. Model fit was further assessed using McFadden's, Cox and Snell's, and Nagelkerke's pseudo-(R^2) indices.

Reliability of Coding

The coding scheme was applied consistently across all three datasets. Intra-coder and inter-coder reliability were estimated. The researcher randomly selected 60 emails (20 emails from

each category of native, non-native, and AI-generated emails) and re-coded the emails for t-units and progression types after four weeks. Cohen's Kappa was estimated. The resulting Cohen's kappa coefficients were $\kappa = .94$ for T-unit identification and $\kappa = .91$ for topical progression classification, indicating very high agreement between the two coding occasions. In addition, a colleague who was trained in coding was asked to code the same subset of data. Cohen's kappa coefficients were $\kappa = .89$ for T-unit identification and $\kappa = .84$ for topical progression classification, indicating substantial agreement according to the interpretation proposed by [Landis and Koch \(1977\)](#). Disagreements between coders were discussed until consensus was achieved.

Findings

The topical progression patterns used by native speakers, non-native speakers, and AI in business emails are examined. Figure 1 compares the proportion of topical progressions in business emails written by native speakers, non-native speakers, and AI. It indicates that Sequential progression, followed by Parallel progression, were the dominant types of progression used in the three groups. On the other hand, Sequential progression was predominantly used by AI. Parallel progression was predominantly used by native speakers and AI. Extended Sequential progression was predominantly used by AI, and Extended Parallel progression was predominantly used by native speakers. To examine if the difference between them is statistically significant, Poisson generalized linear models were run using R software.




Figure 1. The Proportion of Different Topic Progressions Used in Business Emails Written by Native Speakers, Non-Native Speakers, and AI

Are there statistically significant differences among native, non-native, and AI-generated business emails in terms of the frequencies of the four topical structure types after controlling for text length?

Research Question 1 asked whether native, non-native, and AI-generated texts differed in their frequencies of the four topical structure types in the present corpus after controlling for text length. To address this question, separate Poisson generalized linear models were fitted for each topical structure type with the logarithm of the number of t-units included as an offset. Native texts served as the reference category, Incidence Rate Ratios (IRRs) were used to quantify the magnitude of group differences, and Holm-adjusted p-values were used to account for multiple comparisons. As shown in Table 2, the use of Parallel topical structures differed significantly only between native and non-native texts. AI-generated texts did not differ significantly from native texts ($\beta = -0.021$, IRR = 0.98, 95% CI [0.77, 1.25], Holm-adjusted $p = 1.000$). In contrast, non-native texts exhibited a significantly lower rate of Parallel topical structures than native texts ($\beta = -0.681$, IRR = 0.51, 95% CI [0.36, 0.70], Holm-adjusted $p < .001$). This indicates that the frequency of Parallel topical structures used by non-native writers was 49% lower than that of native writers after controlling for text length.

For Sequential topical structures, neither AI-generated texts ($\beta = 0.161$, IRR = 1.17, 95% CI [0.99, 1.40], Holm-adjusted $p = .266$) nor non-native texts ($\beta = 0.030$, IRR = 1.03, 95% CI [0.84, 1.25], Holm-adjusted $p = 1.000$) differed significantly from native texts. Therefore, Sequential topical progression occurred at comparable rates across the three text types.

The analysis of Extended Parallel topical structures revealed significant differences for both comparison groups. AI-generated texts produced substantially fewer Extended Parallel structures than Native texts ($\beta = -1.726$, IRR = 0.18, 95% CI [0.07, 0.40], Holm-adjusted $p < .001$), corresponding to an approximately 82% lower rate. Likewise, non-native texts also demonstrated significantly lower rates of Extended Parallel structures than Native texts ($\beta = -1.169$, IRR = 0.31, 95% CI [0.13, 0.66], Holm-adjusted $p = .026$), representing an approximately 69% reduction in the expected rate.

A different pattern emerged for Extended Sequential topical structures. AI-generated texts produced Extended Sequential structures at a significantly higher rate than native texts ($\beta = 1.048$, IRR = 2.85, 95% CI [1.76, 4.84], Holm-adjusted $p < .001$), indicating that AI-generated texts were expected to contain nearly 2.9 times as many extended sequential structures as native texts. However, the difference between non-native and native texts was

not statistically significant ($\beta = 0.156$, IRR = 1.17, 95% CI [0.61, 2.23], Holm-adjusted $p = 1.000$).

The results demonstrate that the three groups in the present corpus differed selectively rather than uniformly in their topical progression patterns. Native texts were characterized by significantly greater use of Extended Parallel topical progression than both AI-generated and non-native texts, whereas AI-generated texts showed a substantially greater use of Extended Sequential progression. Sequential topical structures occurred at similar rates across all three groups, and Parallel topical structures distinguished only native and non-native texts. These findings suggest that differences in discourse organization among the three text types are concentrated in specific topical progression patterns rather than being evident across all topical structure types.

Table 2. Poisson Generalized Linear Models Predicting Topical Structure Frequencies
(Native Texts as the Reference Category)

Topical Structure	Comparison	B	SE	Z	IRR	95% CI	Holm-adjusted p
Parallel	AI-generated vs. Native	-0.021	0.125	-0.171	0.98	[0.77, 1.25]	1.000
	Non-Native vs. Native	-0.681	0.172	-3.948	0.51	[0.36, 0.70]	< .001
Sequential	AI-generated vs. Native	0.161	0.088	1.835	1.17	[0.99, 1.40]	.266
	Non-Native vs. Native	0.030	0.101	0.295	1.03	[0.84, 1.25]	1.000
Extended parallel	AI-generated vs. Native	-1.726	0.446	-3.870	0.18	[0.07, 0.40]	< .001
	Non-Native vs. Native	-1.169	0.419	-2.794	0.31	[0.13, 0.66]	.026
Extended sequential	AI-generated vs. Native	1.048	0.257	4.076	2.85	[1.76, 4.84]	< .001
	Non-Native vs. Native	0.156	0.330	0.474	1.17	[0.61, 2.23]	1.000

Note. β = regression coefficient (log-rate); SE = standard error; z = Wald statistic; IRR = incidence rate ratio; CI = confidence interval

Can topical progression patterns predict whether a business email is written by a native speaker, a non-native speaker, or AI?

To determine whether topical structure frequencies in the present corpus could predict text type (research question 2), a multinomial logistic regression analysis was conducted. Table 3 presents the overall fit of the multinomial logistic regression model. The likelihood ratio test showed that the full model significantly improved the prediction of text type compared with the intercept-only model, $\chi^2(8) = 223.03$, $p < .001$. The pseudo- R^2 statistics further indicated that the model had substantial explanatory power (McFadden $R^2 = .564$; Cox & Snell $R^2 = .710$; Nagelkerke $R^2 = .799$). In addition, the model achieved an overall classification accuracy of 85.6%, indicating that topical structure frequencies provided strong predictive information for distinguishing among native, non-native, and AI-generated texts.

Table 3. Overall Fit of the Multinomial Logistic Regression Model

Statistics	Value
Likelihood ratio χ^2	223.03
Df	8
P	< .001
McFadden R^2	.564
Cox & Snell R^2	.710
Nagelkerke R^2	.799

The regression coefficients are presented in Table 4. Relative to native texts, AI-generated texts were significantly associated with higher frequencies of Parallel ($\beta = 9.106$, $p = .015$), Sequential ($\beta = 9.381$, $p = .009$), and particularly Extended Sequential topical structures ($\beta = 24.619$, $p < .001$). In contrast, Extended Parallel topical structures significantly reduced the likelihood that a text would be classified as AI-generated ($\beta = -12.121$, $p = .019$).

For the comparison between non-native and native texts, Parallel ($\beta = -17.404$, $p < .001$), Sequential ($\beta = -13.964$, $p < .001$), and Extended Parallel ($\beta = -19.886$, $p < .001$) topical progressions were significant predictors. However, Extended Sequential topical structures did not significantly distinguish non-native from native texts ($\beta = -2.618$, $p = .577$).

Table 4 indicates that not all topical structure types contributed equally to text classification. Extended Sequential structures were the strongest indicator of AI-generated

writing, whereas Extended Parallel structures showed the opposite pattern, making AI-generated texts less likely than native texts. In contrast, non-native texts were characterized by significantly lower rates of Parallel, Sequential, and Extended Parallel structures relative to native texts, while Extended Sequential structures did not differentiate non-native and native writing. These findings suggest that the distribution of topical structures varies systematically across text types and provides meaningful linguistic cues for distinguishing human and AI-generated writing.

Table 4. Multinomial Logistic Regression Predicting Text Type (Reference Category = Native)

Comparison	Predictor	B	SE	Z	P	OR	95% CI
AI-generated	Parallel	9.106	3.729	2.442	.015	9005.64	6.03–13,45
	Sequential	9.381	3.595	2.610	.009	11862.62	10.34–13,61
	Extended Parallel	–12.121	5.166	–2.346	.019	0.000005	0.000000000218–0.13
	Extended Sequential	24.619	4.784	5.146	< .001	4.92×10^{10}	4.16×10^6 – 5.81×10^{14}
Non-Native	Parallel	–17.404	3.480	–5.000	< .001	2.76×10^{-8}	3.01×10^{-11} – 2.54×10^{-5}
	Sequential	–13.964	3.564	–3.918	< .001	8.62×10^{-7}	7.97×10^{-10} –0.000932
	Extended Parallel	–19.886	5.088	–3.908	< .001	2.31×10^{-9}	1.08×10^{-13} –0.000050
	Extended Sequential	–2.618	4.693	–0.558	.577	0.073	0.000007–721.24

As the confusion matrix (Table 5) shows, the multinomial logistic regression correctly classified 154 of the 180 texts (85.6%). AI-generated and non-native texts were identified with particularly high accuracy, with no AI-generated texts being classified as non-native and no non-native texts being classified as AI-generated. Most classification errors occurred between native texts and the other two groups, suggesting that native texts shared some topical structure characteristics with both AI-generated and non-native writing. Nevertheless, the overall classification accuracy indicates that topical structure frequencies provide strong discriminatory information for identifying text type.

Table 5. Classification Accuracy of the Multinomial Logistic Regression Model

Actual Text	Predicted Native	Predicted AI-generated	Predicted Non-Native
Native	47	6	7
AI-generated	7	53	0
Non-Native	6	0	54
Overall classification accuracy = 85.6%			

Discussion

Are there statistically significant differences among native, non-native, and AI-generated business emails in terms of the frequencies of the four topical structure types after controlling for text length?

The findings showed that the three groups studied in the corpus of the present study did not differ uniformly across all topical progression patterns. Although significant differences were observed in Parallel, Extended Parallel, and Extended Sequential progression, Sequential progression occurred comparably in the three groups. Thus, Sequential progression is a discourse strategy shared by both human writers and AI. It is one of the basic mechanisms for maintaining coherence since it presents ideas that are developed from the previous clause (Chang, 2023). This interpretation is consistent with research studies reporting sequential progression as the predominant or one of the most frequently occurring topical progressions in L1 or L2 writing contexts (Herriman, 2011; Schneider & Connor, 1990; Shabana, 2018; Zhang & Lee, 2019). One possible explanation is that Sequential progression is closely aligned with the general communicative need to develop information progressively in business correspondence.

Another notable finding was the greater use of Extended Parallel progression in emails written by native speakers. Extended Parallel progression, in which the topic of a sentence returns to a topic established several sentences earlier, represents a cognitively demanding discourse strategy. It requires the writer to maintain a mental representation of earlier discourse topics and to reactivate them purposefully, thereby creating higher-quality texts; it facilitates making semantic ties in texts (Witte & Faigley, 1981). The greater frequency of Extended Parallel progression in the native-speaker group may therefore indicate a greater tendency to reactivate established discourse topics when developing business messages. This progression helps readers return to the important topics in the text and elaborate on them (Witte, 1983).

The significantly lower use of Extended Parallel progression by non-native speakers is consistent with [Uysal \(2012\)](#), who demonstrated that L2 writers fail to use discourse strategies such as rhetorical presentations. In addition, it can be justified by [Kellogg et al. \(2013\)](#), who argued that novice writers devote substantial cognitive resources to lexical retrieval and grammatical encoding, leaving fewer attentional resources available for global discourse planning. Consequently, L2 texts often achieve adequate local coherence while exhibiting weaker long-distance thematic integration. The present findings also reinforce [Incelli's \(2013\)](#) argument that discourse competence develops more gradually than grammatical competence.

A contrasting pattern emerged for Extended Sequential progression. AI-generated emails used this progression at nearly three times the rate observed in native texts, whereas non-native writers did not differ significantly from native speakers. In Extended Sequential progression, later topics are derived from information introduced in the rheme of earlier clauses ([Simpson, 2000](#)). This progression may reflect the large language model's (LLM) generation process: recently introduced lexical and semantic material strongly conditions subsequent output, encouraging repeated expansion of prior content ([Elkhatat, 2023](#)). This pattern is also in line with research showing that AI-generated texts tend to have higher rates of repetition and lower lexical diversity ([Gao et al., 2023](#)).

Another finding concerns Parallel progression. Only native and non-native writers differed significantly; non-native writers produced lower Parallel progression. Parallel progression enables writers to maintain a stable topical focus while elaborating multiple aspects of the same topic ([Lautamatti, 1987](#)). Previous TSA research associated greater use of Parallel progression with higher coherence ([Connor & Farmer, 1990](#); [Schneider & Connor, 1990](#)). The reduced frequency observed among non-native writers therefore suggests difficulty maintaining thematic stability through business communication. Research indicates that EFL learners frequently prioritize introducing new information over sustaining an existing topic, resulting in texts that are coherent locally but less integrated globally ([Flores & Yin, 2015](#); [Liao et al., 2023](#)).

By contrast, AI-generated texts did not differ significantly from native texts in their use of Parallel progression. This suggests that LLMs can replicate the surface-level topic continuity characteristic of native expert writing. This finding is consistent with research indicating that LLMs can produce locally coherent text that maintains referential continuity across adjacent sentences ([Gao et al., 2023](#)). However, as the findings for Extended Parallel

and Extended Sequential progressions demonstrate, this local coherence does not extend uniformly to more complex, long-range topical organization.

Can topical progression patterns predict whether a business email is written by a native speaker, a non-native speaker, or AI?

The analyses identified Extended Sequential progression as the strongest positive predictor of AI authorship, while Extended Parallel progression was the strongest negative predictor in the corpus investigated in this study. This finding is consistent with the group-level comparisons in RQ1. This pattern reveals a distinctive discourse pattern for AI-generated business emails: they are characterized by a particular topical progression in which ideas are elaborated in long forward-moving chains (Extended Sequential) rather than being revisited through topic returns (Extended Parallel). This profile constitutes a form of AI discourse fingerprint detectable by TSA, extending the concept of AI fingerprints previously identified at the lexical and syntactic levels (Casal & Kessler, 2023; Gao et al., 2023; Roshid et al., 2018).

For non-native texts, Parallel, Sequential, and Extended Parallel progressions were all significant negative predictors compared with native-speaker texts. This finding indicates that non-native writers in this corpus used all three of these topical structure types at significantly lower rates than native writers; this pattern suggests that non-native business emails are characterized by a broadly reduced deployment of topical progression strategies, potentially reflecting shorter texts, simpler discourse organization, or avoidance of complex rhetorical structures (Shabana, 2018; Uysal, 2012). Importantly, Extended Sequential progression did not significantly distinguish non-native from native texts.

These findings extend prior research on native-non-native differences in communication. Studies of business email and letter writing have demonstrated that non-native writers differ from native writers not only in grammatical accuracy but in the deployment of rhetorical and discourse-level structures (Flowerdew & Wan, 2010; Incelli, 2013; Mehrpour & Mehrzad, 2013; Park et al., 2021; Van Herck et al., 2022). The present study adds topical structure distribution to this list of differentiating features and demonstrates that these differences are strong enough to support text classification with high accuracy.

Finally, the confusion matrix revealed a particularly notable result: no AI-generated text was classified as non-native, and no non-native text was classified as AI-generated. This exclusivity indicates that, despite both groups differing from native writing, AI and non-

native texts have distinct topical structure. This finding challenges the assumption that AI-generated text resembles non-native writing, a concern raised by [Liang et al. \(2023\)](#), who documented that AI detection tools disproportionately flag non-native writing as AI-generated. The present study suggests that, at the level of topical structure, AI and non-native writing are distinct, with AI overproducing Extended Sequential chains and non-native underproducing all major topical structure types.

Conclusions

This study investigated differences in topical progression patterns among business emails written by native speakers, non-native speakers, and AI. Specifically, it examined whether the three groups differed in the frequencies of four topical progression patterns after controlling for text length and whether these discourse features could predict text authorship. The findings showed that discourse organization varied systematically across the three text types. Rather than differing uniformly across all topical progression patterns, the groups exhibited distinctive discourse patterns. Native writers used significantly more Extended Parallel progression than non-native writers and AI, reflecting recursive thematic development. In contrast, AI-generated texts relied heavily on Extended Sequential progression. Moreover, Sequential progression occurred at similar rates across all groups, while Parallel progression differentiated only native and non-native texts. Furthermore, multinomial logistic regression showed that topical structure frequencies predicted text type with high accuracy (85.6%), highlighting the effectiveness of discourse-level features in distinguishing native, non-native, and AI-generated business emails.

The present findings contribute to the theoretical development of TSA in several ways. First, they demonstrate that the TSA framework, originally developed for academic writing ([Lautamatti, 1978](#); [Witte, 1983](#)), retains its discriminatory power when applied to the professional genre of business email and AI-generated writing. Second, the differential use of Extended Parallel versus Extended Sequential progression shows the difference between human expert writing and AI-generated texts. Third, the finding that non-native and AI texts differed from native texts in distinct ways suggests that the mechanisms underlying their departures from native discourse norms are different: non-native writers are constrained by L2 proficiency limitations and avoidance strategies, while AI is constrained by features that privilege local, forward-moving elaboration over global, retrospective topic management.

The findings also have implications for the teaching of professional writing to non-native business communicators. Given that non-native writers in this study systematically

underused Parallel, Extended Parallel, and Sequential topical progressions relative to native writers, explicit instruction in topical structure may be beneficial in business English programs. Instruction that draws learners' attention to the importance of topic return (Extended Parallel) and sustained topic focus (Parallel) may help non-native writers produce more coherent and effective business emails.

Despite its contributions, the study has several limitations. First, the corpus consisted exclusively of business emails. Thus, the findings may not generalize to other forms of writing. Future studies should examine topical progression across a wider range of professional and academic genres. Second, the findings concerning AI-generated writing should also be interpreted within the specific boundaries of the present corpus. The AI-generated emails were produced using a single LLM (ChatGPT-5) and a specific prompting procedure. Consequently, the results demonstrate that the topical progression patterns examined in this study provide discriminating information among the three groups represented in the present dataset. Different LLMs may produce different discourse patterns, and changes in prompts, model versions, generation settings, or communicative contexts may affect topical organization. Future research should test the robustness of TSA across outputs generated by multiple LLMs, different prompting conditions, model versions, and larger corpora. Third, another limitation concerns the unequal total number of words across the three groups. To reduce the influence of text length on the statistical comparisons, the Poisson models incorporated the logarithm of the number of T-units as an offset, and the predictors in the multinomial logistic regression were expressed as frequencies per T-unit. Nevertheless, normalization does not make the underlying corpora equivalent in size. Therefore, the findings should be interpreted as evidence from the present corpus rather than as definitive evidence about all native, non-native, or AI-generated business emails. Future research should replicate the analysis using more closely balanced corpora, ideally with comparable numbers of words and T-units across groups.

Declaration of Conflicting Interest

There is no conflict of interest in conducting this study.

References

- AlAfnan, M. (2025). Technical report writing efficiency using AI-powered tools: Opportunities, challenges, and future directions. *Journal of Artificial Intelligence and Technology*, 5, 270-277. <https://doi.org/10.37965/jait.2025.0729>

- Cardon, P., Fleischmann, C., Logemann, M., Heidewald, J., Aritz, J., & Swartz, S. (2023). Competencies needed by business professionals in the AI age: Character and communication lead the way. *Business and Professional Communication Quarterly*, 87(2), 223–246. <https://doi.org/10.1177/23294906231208166>
- Casal, J. E., & Kessler, M. (2023). Can linguists distinguish between ChatGPT/AI and human writing? A study of research ethics and academic publishing. *Research Methods in Applied Linguistics*, 2(3), 100068. <https://doi.org/10.1016/j.rmal.2023.100068>
- Chang, P. (2023). Reading research genre: The impact of thematic progression. *REL C Journal*, 54(1), 129–148. <https://doi.org/10.1177/00336882211013613>
- Connor, U., & Farmer, M. (1990). The teaching of topical structure analysis as a revision strategy for ESL writers. In B. Kroll (Ed.), *Second language writing: Research insights for the classroom* (pp. 126–139). Cambridge University Press. <https://doi.org/10.1017/CBO9781139524551.013>
- Elkhatat, A. M. (2023). Evaluating the authenticity of ChatGPT responses: A study on text-matching capabilities. *International Journal for Educational Integrity*, 19(1), 15–31. <https://doi.org/10.1007/s40979-023-00137-0>
- Flores, E. R., & Yin, K. (2015). Topical structure analysis as an assessment tool in student academic writing. *3L: The Southeast Asian Journal of English Language Studies*, 21(1), 103–115. <https://doi.org/10.17576/31-2015-2101-10>
- Flowerdew, J., & Wan, A. (2010). The linguistic and the contextual in applied genre analysis: The case of the company audit report. *English for Specific Purposes*, 29(2), 78–93. <https://doi.org/10.1016/j.esp.2009.07.001>
- Gao, C. A., Howard, F. M., Markov, N. S., Dyer, E. C., Ramesh, S., Luo, Y., & Pearson, A. T. (2023). Comparing scientific abstracts generated by ChatGPT to real abstracts with detectors and blinded human reviewers. *Npj Digital Medicine*, 6(1), Article 75. <https://doi.org/10.1038/s41746-023-00819-6>
- Getchell, K. M., Carradini, S., Cardon, P. W., Fleischmann, C., Ma, H., Aritz, J., & Stapp, J. (2022). Artificial intelligence in business communication: The changing landscape of research and teaching. *Business and Professional Communication Quarterly*, 85(1), 7–33. <https://doi.org/10.1177/23294906221074311>
- Haleem, A., Javaid, M., & Singh, R. P. (2023). An era of ChatGPT as a significant futuristic support tool: A study on features, abilities, and challenges. *Bench Council Transactions on Benchmarks, Standards and Evaluations*, 2(4), 100089. <https://doi.org/10.1016/j.tbench.2023.100089>

- Halliday, M. A. K. (2009). Methods – techniques – problems. In M. A. K. Halliday, & J. J. Webster (Eds.), *Continuum companion to systemic functional linguistics* (pp. 59–86). Continuum
- He, J., Weerkamp, W., Larson, M. & de Rijke, M. (2009). An effective coherence measure to determine topical consistency in user-generated content. *International Journal on Document Analysis and Recognition*, 12, 185-203. <https://doi.org/10.1007/s10032-009-0089-5>
- Herbold, S., Hautli-Janisz, A., Heuer, U., Kikteva, Z., & Trautsch, A. (2023). A large-scale comparison of human-written versus ChatGPT-generated essays. *Scientific Reports*, 13(1), Article 18617. <https://doi.org/10.1038/s41598-023-45644-9>
- Herriman, J. (2011). Themes and theme progression in Swedish advanced learners' writing in English. *Nordic Journal of English Studies*, 10(1), 1–28. <https://doi.org/10.35360/njes.240>
- Huddleston, R. (1984). *Introduction to the grammar of English*. Cambridge University Press.
- Incelli, E. (2013). Managing discourse in intercultural business email interactions: A case study of a British and Italian business transaction. *Journal of Multilingual and Multicultural Development*, 34, 515-532. <https://doi.org/10.1080/01434632.2013.807270>
- Kellogg, R. T., Whiteford, A., P., Turner, C., E., Cahill, M. & Mertens (2013). Working memory in written composition: A progress report. *Journal of Writing Research*, 5(2), 159-190. <https://doi.org/10.17239/jowr-2013.05.02.1>
- Landis, J. R., & Koch, G. G. (1977). The measurement of observer agreement for categorical data. *Biometrics*, 33(1), 159-174. <https://doi.org/10.2307/2529310>
- Lautamatti, L. (1987). "Observations on the development of the topic of simplified discourse. In U. Connor & R. B. Kaplan (Eds.), *Writing across languages: Analysis of L2 text* (pp. 87-114). Addison-Wesley
- Liang, W., Yuksekgonul, M., Mao, Y., Wu, E., & Zou, J. (2023). GPT detectors are biased against non-native English writers. *Patterns*, 4(7), 100779. <https://doi.org/10.1016/j.patter.2023.100779>
- Liao, W., Liu, Z., Dai, H., Xu, S., Wu, Z., Zhang, Y., Huang, X., Zhu, D., Cai, H., Li, Q., Liu, T., & Li, X. (2023). Differentiating ChatGPT-generated and human-written medical texts: Quantitative study. *JMIR Medical Education*, 9, e48904. <https://doi.org/10.2196/48904>

- Liu, P. & Liu, H. (2023). Interpersonal strategies in international business emails: The intercultural pragmatics perspective. *Intercultural Pragmatics*, 20, 557-579. <https://doi.org/10.1515/ip-2023-5004>
- Mehrpour, S., & Mehrzad, M. (2013). A comparative genre analysis of English business e-mails written by Iranians and native English speakers. *Theory and Practice in Language Studies*, 3(12), 2250-2261. <https://doi.org/10.4304/tpls.3.12.2250-2261>
- Olszewski, R., Brzeziński, J., Watros, K., & Rysz, J. (2026). Quantifying readability in chatbot-generated medical texts using classical linguistic indices: A review. *Applied Sciences*, 16(3), 1423. <https://doi.org/10.3390/app16031423>
- Papakonstantinidis, S. (2024). Embrace or resist? Drivers of artificial intelligence writing software adoption in academic and non-academic contexts. *Contemporary Educational Technology*, 16(2), ep495. <https://doi.org/10.30935/cedtech/14250>
- Park, S., Jeon, J., & Shim, E. (2021). Exploring request emails in English for business purposes: A move analysis. *English for Specific Purposes*, 63, 137–150. <https://doi.org/10.1016/j.esp.2021.03.006>
- Qu, G. (2024). A comparative study of politeness and hedges used in corporate emails written by business English majors, AI, and native speakers. *Frontiers in Educational Research*, 7(12), 118-125. <https://doi.org/10.25236/fer.2024.071218>
- Quinto, E. (2015). Physical and topical structures of manpower discourse: A contrastive rhetoric analysis in Southeast Asia. *GEMA Online Journal of Language Studies*, 15, 91-110. <https://doi.org/10.17576/GEMA-2015-1501-06>
- Reviriego, P., Conde, J., Merino-Gómez, E., Martínez, G., & Hernández, J. A. (2024). Playing with words: Comparing the vocabulary and lexical diversity of ChatGPT and humans. *Machine Learning with Applications*, 18, 100602. <https://doi.org/10.1016/j.mlwa.2024.100602>
- Roshid, M. M., Webb, S., & Chowdhury, R. (2018). English as a business lingua franca: A discursive analysis of business e-mails. *International Journal of Business Communication*, 59(1), 83–103. <https://doi.org/10.1177/2329488418808040>
- Saeed, A., Everatt, J., Sadeghi, A., & Munir, A. (2022). Cognitive predictors of coherence in adult ESL learners' writing. *Journal of Language and Education*, 8(3), 106-118. <https://jle.hse.ru/article/view/12935>
- Schneider, M., & Connor, U. (1990). Analyzing topical structure in ESL essays. *Studies in Second Language Acquisition*, 12(4), 411–427. <https://doi.org/10.1017/s0272263100009505>

- Shabana, N. O. (2018). Topical structure analysis: Assessing first-year Egyptian university students' internal coherence of their EFL writing. In A. H. Ahmed & H. Abouabdelkader (Eds.), *Assessing EFL writing in the 21st century Arab world* (pp. 53–78). Palgrave Macmillan. https://doi.org/10.1007/978-3-319-64104-1_3
- Simpson, J. M. (2000). Topical structure analysis of academic paragraphs in English and Spanish. *Journal of Second Language Writing*, 9(3), 293–309. [https://doi.org/10.1016/s1060-3743\(00\)00029-1](https://doi.org/10.1016/s1060-3743(00)00029-1)
- Sujinpram, N. & Wannaruk, A. (2026). From inbox to insight: Materials design for global business email communication. *rEFLections*, 33(1), 169-193. <https://doi.org/10.61508/refl.v33i1.288455>
- Swales, J. M. (1990). *Genre analysis: English in academic and research settings*. Cambridge University Press.
- Tsai, S. (2025). Online EFL business writing with GenAI-generated templates: Students' performance and perceptions. *Australasian Journal of Educational Technology*, 41(6), 82-97. <https://doi.org/10.14742/ajet.10743>
- Uysal, H. H. (2012). Argumentation across L1 and L2 writing: Exploring the transfer and non-transfer of discourse features. *VIAL*, 9, 133-159.
- Van Herck, R., Decock, S., & Fastrich, B. (2022). A unique blend of interpersonal and transactional strategies in English email responses to customer complaints in a B2C setting: A move analysis. *English for Specific Purposes*, 65, 30–48. <https://doi.org/10.1016/j.esp.2021.08.001>
- Wang, J. (2024). Improving ChatGPT's competency in generating effective business communication messages: integrating rhetorical genre analysis into prompting techniques. *Journal of Technical Writing and Communication*, 54, 369-395. <https://doi.org/10.1177/00472816241260033>
- Wilson, W., & Rose, H. (2025). A genre, scoring, and authorship analysis of AI-generated and human-written refusal emails. *Business and Professional Communication Quarterly*. Advance online publication. <https://doi.org/10.1177/23294906251322890>
- Witte, S. P. (1983). Topical structure and writing quality: Some possible text-based explanations of readers' judgments of student writing. *Visible Language*, 17(2), 177–205.
- Witte, S. P., & Faigley, L. (1981). Coherence, cohesion, and writing quality. *College Composition and Communication*, 32(2), 189–204. <https://doi.org/10.58680/cc198115912>

- Yang, S., Chen, S., Zhu, H., Lin, J., & Wang, X. (2024). A comparative study of thematic choices and thematic progression patterns in human-written and AI-generated texts. *System*, 126, 103494. <https://doi.org/10.1016/j.system.2024.103494>
- Zhang, Z., & Lee, B. (2019). Thematic progression patterns in English abstracts of doctoral dissertations by EFL students. *Korean Journal of English Language and Linguistics*, 19(4), 668–687. <https://doi.org/10.15738/kjell.19.4.201912.668>

Appendix A

Schneider and Connor's guideline to examine topic progression patterns (1990, p. 427):

A) Parallel Progression:

1. Any topic that exactly repeats, is a pronominal form, or is a synonym of the immediately preceding sentence topic.
2. Any topic that is a singular or plural form of the immediately preceding sentence topic.
3. Any topic that is an affirmative or negative form of the immediately preceding sentence topic (e.g., artists, no artists).

B) Sequential Progression:

1. Any topic that is different from the immediately preceding topic, that is, not points (1)-(3) mentioned in the previous part.
2. Any topic in which there is a qualify er that so limits or further specifies a noun phrase that it refers to a different referent (e.g., a nation; a very small, multi-racial nation, referring to two different nations).
3. Any topic that is a derivation of an immediately preceding topic (science, scientists).
4. Any topic that is related to the immediately preceding topic by a part-whole relationship (e.g., these groups, housewives, children, and old people).
5. Any topic that repeats a part but not all of an immediately preceding sentence topic (e.g., science and art, science, art).

C) Extended Parallel Progression:

Any topic which is interrupted by at least one sequential topic before it returns to a previous topic.

Simpson (2000) added the fourth type of topic progression, Extended Sequential progression, in which a unit in the rheme of an earlier t-unit became the topic of a later t-unit after an intervening sequence.

Appendix B

Prompt for AI-generated business email

Write a professional business email in English based on the communicative context provided below.

The email should:

- address the communicative purpose described in the context;
- be appropriate for a professional business setting;
- communicate the necessary information clearly and completely;

Do not explain your choices or provide any analysis. Write only the email.

Communicative context: [INSERT CONTEXT HERE]