

ROPE: Reading Order Equivariant Positional Encoding for Graph-based Document Information Extraction

Chen-Yu Lee[†], Chun-Liang Li[†], Chu Wang^{§*}, Renshen Wang[‡],
Yasuhisa Fujii[‡], Siyang Qin[‡], Ashok Papat[‡], Tomas Pfister[†]

[†]Google Cloud AI, [§]McGill University, [‡]Google Research

^{†,‡}{chenyulee, chunliang, rewang, yasuhisaf, qinb, popat, tpfister}@google.com

[§]chu.wang@mail.mcgill.ca

Abstract

Natural reading orders of words are crucial for information extraction from form-like documents. Despite recent advances in Graph Convolutional Networks (GCNs) on modeling spatial layout patterns of documents, they have limited ability to capture reading orders of given word-level node representations in a graph. We propose Reading Order Equivariant Positional Encoding (ROPE), a new positional encoding technique designed to apprehend the sequential presentation of words in documents. ROPE generates unique reading order codes for neighboring words relative to the target word given a word-level graph connectivity. We study two fundamental document entity extraction tasks including word labeling and word grouping on the public FUNSD dataset and a large-scale payment dataset. We show that ROPE consistently improves existing GCNs with a margin up to 8.4% F1-score.

1 Introduction

Key information extraction from form-like documents is one of the fundamental tasks of document understanding that has many real-world applications. However, the major challenge of solving the task lies in modeling various template layouts and formats of documents. For example, a single document may contain multiple columns, tables, and non-aligned blocks of texts (e.g. Figure 1).

The task has been studied from rule-based models (Lebourgeois et al., 1992) to learning-based approaches (Palm et al., 2017; Tata et al., 2021). Inspired by the success of sequence tagging in NLP (Sutskever et al., 2014; Vaswani et al., 2017; Devlin et al., 2019), a natural extension is applying these methods on linearly serialized 2D documents (Palm et al., 2017; Aggarwal et al., 2020).

* Work done while an intern at Google Research.




Figure 1: Illustration of the proposed Reading Order Equivariant Positional Encoding (ROPE). **Top:** a portion of a form document with the original word reading order. **Bottom:** given a graph connectivity, ROPE generates equivariant reading order codes with respect to the target word (in this case the date “3/18/97”).

Nevertheless, scattered columns, tables, and text blocks in documents make the serialization extremely difficult, largely limiting the performance of sequence models. Katti et al. (2018); Zhao et al. (2019) explore to directly work on 2D document space using grid-like convolutional models to better preserve spatial context during learning, but the performance is restrictive to the resolution of the grids. Recently, Qian et al. (2019); Davis et al. (2019); Liu et al. (2019) propose to represent documents using graphs, where nodes define word tokens and edges describe the spatial patterns of words. Yu et al. (2020) show state-of-the-art performance of Graph Convolutional Networks (GCNs) (Duvenaud et al., 2015) on document understanding.

Although GCNs capture the relative spatial relationships between words through edges, the specific word ordering information is lost during the graph aggregation operation, in the similar way to the average pooling in Convolutional Neural Networks (CNNs). However, we believe reading orders are strong prior to comprehending languages. In this work, we propose a simple yet effective Reading Order Equivariant Positional Encoding (ROPE) that embeds the relative reading order context into

graphs, bridging the gap between sequence and graph models for robust document understanding. Specifically, for every word in a constructed graph, ROPE generates unique reading order codes for its neighboring words based on the graph connectivity. The codes are then fed into GCNs with self-attention aggregation functions for effective relative reading order encoding. We study two fundamental entity extraction tasks including word labeling and word grouping on the public FUNSD dataset and a large-scale payment dataset. We observe that by explicitly encoding relative reading orders, ROPE brings the same or higher performance improvement compared to spatial relationship features in existing GCNs in parallel.

2 Other Related Work

Attention models show state-of-the-art results in graph learning (Veličković et al., 2018) and NLP benchmarks (Vaswani et al., 2017). As attention models with positional encodings are proven to be universal approximators of sequence-to-sequence functions (Yun et al., 2020), encoding positions or ordering is an important research topic. For sequence, learned positional embeddings (Gehring et al., 2017; Devlin et al., 2019; Shaw et al., 2018), sinusoidal functions and its extensions (Liu et al., 2020) have been studied. Beyond that, positional encodings are explored in graphs (You et al., 2019), 2D images (Parmar et al., 2018) and 3D structures (Fuchs et al., 2020). Lastly, graph modeling is also applied to other document understanding tasks, including document classification (Yao et al., 2019) and summarization (Yasunaga et al., 2017).

3 Method

We follow recent advances in using GCNs for document information extraction that relax any serialization assumptions by sequence modeling. GCNs take inputs (word tokens in this case) of arbitrary numbers, sizes, shapes and locations, and encode the underlying spatial layout patterns of documents through direct message passing and gradient updates between input embedding in the 2D space.

Node definition. Given a document D with N tokens denoted by $T = \{t_1, t_2, \dots, t_N\}$, we refer t_i to the i -th token in a linearly serialized text sequence returned by the Optical Character Recognition (OCR) engine. The OCR engine generates the bounding box sizes and locations for all tokens, as

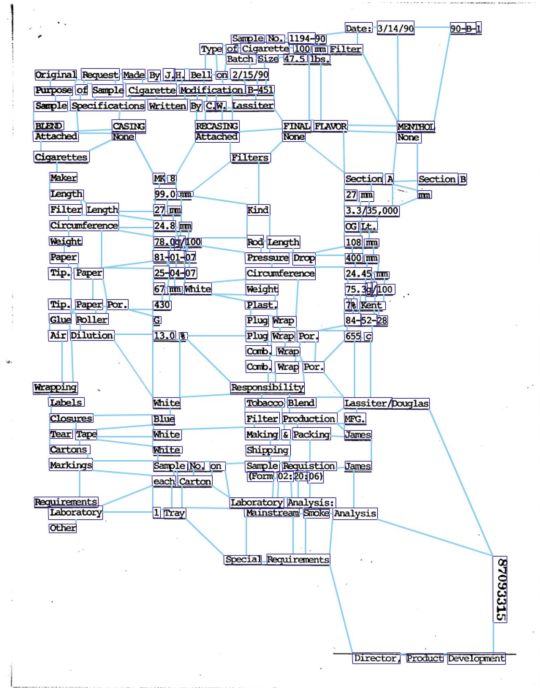


Figure 2: Sample of a β -skeleton graph of a document of FUNSD.

well as the text within each box. We define node input representation for all tokens T as vertices $V = \{v_1, v_2, \dots, v_N\}$, where v_i concatenates quantifiable attributes available for t_i . In our design, we use two common input modalities: (a) word embeddings from an off-the-shelf pre-trained BERT model (Devlin et al., 2019), and (b) spatial embeddings from normalized bounding box heights, widths, and Cartesian coordinate values of four corners.

Edge definition. While the vertices V represent tokens in a document, the edges characterize the relationship between the vertices. Precisely, we define directional edges for a set of edges E , where each edge e_{ij} connects two vertices v_i and v_j , concatenating quantifiable edge attributes. In our design, we use two input modalities given an edge e_{ij} connecting two vertices: (a) spatial embeddings from horizontal and vertical normalized relative distances between centers, top left corners and bottom right corners of the bounding boxes. It also contains height and width aspect ratios of v_i , v_j , and relative height and width aspect ratios between v_i and v_j . (b) Visual embeddings that utilizes ImageNet pre-trained MobileNetV3 (Howard et al., 2019) to extract visual representations of union bounding boxes containing v_i and v_j . The visual embedding in edge formation picks up visual cues




Figure 3: Implementation of the proposed Reading Order Equivariant Positional Encoding (ROPE). Given a graph connectivity, ROPE iterates through the neighboring word vertices in the original reading order and assigns new ROPE codes (red numbers) to the neighbors, starting from zero. Note that the proposed ROPE codes remain unchanged if the neighbors and the target shift equally in the document with the same relative reading order, therefore being equivariant.

such as colors, fonts, separating symbols or lines between two token bounding boxes (through their union bounding box). We refer to the spatial embedding in (a) as the edge geometric (EdgeGeo) feature used in the experimental section.

Graph construction. Our implementation is based on the β -skeleton graph (Kirkpatrick and Radke, 1985) with $\beta = 1$ for graph construction. By using the “ball-of-sight” strategy, β -skeleton graph offers high connectivity between word vertices for necessary message passing while being much sparser than fully-connected graphs for efficient forward and backward computations (Wang et al., 2021). A β -skeleton graph example can be found in Figure 2, and more can be found in Figure 5 in the Appendix.

Aggregation function. Inspired by the Graph Attention Networks (Veličković et al., 2018) and the Transformers (Vaswani et al., 2017), we use multi-head self-attention module as our GCN aggregation (pooling) function. It calculates the importance of individual message coming from its neighbors to generate the new aggregated output.

3.1 Reading Order Equivariant Positional Encoding (ROPE)

Positional encoding (Gehring et al., 2017) in sequence models is with an assumption that the input is perfectly serialized. However, as illustrated in Figure 1, form-like documents often contain multiple columns or sections. A simple left-to-right and top-to-bottom serialization commonly provided by OCR engines does not provide accurate sequential presentation of words – two consecutive words in the same sentence might have drastically different reading order indexes by naive serialization.

Instead of assigning absolute reading order indexes for the entire document at the beginning, we

propose to encode the relative reading order context of neighboring words w.r.t. the target word based on the given graph connectivity. Figure 3 demonstrates the process of the proposed method: ROPE iterates through the neighboring word vertices in the original reading order and assigns new ROPE codes $p \in \mathbb{N}$ (red numbers) to the neighbors, starting from zero. The generated codes are then appended to the corresponding incoming messages during graph message passing. Hence, ROPE provides a relative reading order context of the neighborhood for order-aware self-attention pooling.

Note that the generated ROPE codes remain unchanged if the neighbors and the target shift equally in the document with the same relative order, therefore being equivariant. Additionally, ROPE provides robust sequential output that is consistent even when the neighborhood crosses multiple columns or sections in a document.

Finally, we also explore sinusoidal encoding matrix (Vaswani et al., 2017) besides the index-based encoding. Our ablation study in Section 4 shows that using both results in the best performance.

4 Experiments

We evaluate how reading order impacts overall performance of graph-based information extraction from form-like documents. We adopt two form understanding tasks as Jaume et al. (2019), including word labeling and word grouping. Word labeling is the task of assigning each word a label from a set of predefined entity categories, realized by node classification. Word grouping is the task of aggregating words that belong to the same entity, realized by edge classification. These two fundamental entity extraction tasks do not rely on perfect entity word groupings provided by the dataset and therefore help decouple the modeling capability provided by the proposed ROPE in practice. These two tasks

also effectively demonstrate the quality of the node embedding and edge embedding of the proposed graph architecture and decouple any performance gain from sophisticated Conditional Random Field (CRF) decoders often used on top of the model.

4.1 Datasets

Payment. We follow Majumder et al. (2020) to prepare a large-scale payment document collection that consists of around 18K single-page payments. The data come from different vendors with different layout templates. For both word labeling and word grouping experiments, we use a 80-20 split of the corpus as the training and test sets.

We use a public OCR service¹ to extract words from the payment documents. The service generates the text of each word with their corresponding 2D bounding box. The word boxes are roughly arranged in an order from left to right and from top to bottom. We then ask human annotators to label the words with 13 semantic entities. Each entity ground truth is described by an entity type and a list of words generated by the OCR engine, resulting in over 3M word-level annotations. Labelers are instructed to label all instances of a field in a document, therefore our GCNs are trained to predict all instances of a field as well.

FUNSD. FUNSD (Jaume et al., 2019) is a public dataset for form understanding in noisy scanned documents, containing a collection of research, marketing, and advertising documents that vary widely in their structure and appearance. The dataset consists of 199 annotated forms with 9,707 entities and 31,485 word-level annotations for 4 entity types: header, question, answer, and other. For both word labeling and word grouping experiments, we use the official 75-25 split for the training and test sets.

4.2 Experimental Setup

All GCN variants used in the experiment have the same architecture: The node update function is a 2-layer Multi-Layer Perceptron (MLP) with 128 hidden nodes. The aggregation function uses a 3-layer multi-head self-attention pooling with 4 heads and 32 as the head size. The number of hops in the GCN is set to 7 for payment dataset and 2 for FUNSD dataset due to the complexity and scale of the former. We use cross-entropy loss for both multi-class word labeling and binary word

¹cloud.google.com/vision

	Types of Positional Encoding (ours)		Word Labeling	Word Grouping		
	EdgeGeo	ROPE	F1	P	R	F1
Payment	✓		60.80	83.64	83.97	83.80
		✓	66.09	84.96	84.93	84.94
	✓	✓	68.17	84.92	86.86	85.88
FUNSD			50.86	82.09	92.21	86.86
	✓		53.16	87.56	87.17	87.37
	✓	✓	51.78	88.90	89.67	89.28
	✓		57.22	88.64	90.03	89.33

Table 1: Different positional encodings for GCNs on information extraction tasks. We observe that the reading order encoding (ROPE) is equally or more important compared to edge geometric feature (EdgeGeo).

	ROPE Encoding Function		Word Labeling	Word Grouping		
	Index	Sinusoidal	F1	P	R	F1
Payment	✓		66.09	84.96	84.93	84.94
		✓	72.41	87.78	85.31	86.53
	✓	✓	70.94	88.49	83.00	85.66
FUNSD			53.16	87.56	87.17	87.37
	✓		55.48	85.95	92.15	88.94
	✓	✓	54.14	88.72	89.51	89.12
	✓		57.22	88.64	90.03	89.33

Table 2: Ablation of positional encoding function used in the proposed ROPE. We observe that either index or sine encoding works better than no positional encoding. Combined works the best.

grouping tasks. We train the models from scratch using Adam optimizer with the batch size of 1. The learning rate is set to 0.0001 with warm-up proportion of 0.01. The training is conducted on 8 Tesla P100 GPUs for approximately 1 day on the largest corpus.

4.3 Results

We train the GCNs from scratch on all datasets. For word labeling we use multi-class node classification F1-scores as the metric and for word grouping we use binary edge classification F1-scores as the metric with the corresponding precision and recall values.

Importance of reading order. Positional encoding mechanisms are the key components to exploiting layout patterns of words – Answer entities are usually next to or below the Question entities. Existing GCN approaches rely on edge geometric (EdgeGeo) features to capture such spatial relationships between words in 2D space. Here we evaluate the importance of the proposed reading order encoding ROPE with various combinations of EdgeGeo over the baseline GCN (Qian et al.,




Figure 4: Sensitivity of ROPE to OCR reading order on Payment. The proposed ROPE codes remain the same if the connected neighboring words and target word shift equally in the document.

2019) as summarized in Table 1. Without any positional encoding, word labeling F1 drops by 13.75 points and word grouping F1 drops by 2.84 points on payment dataset. Then, we pass ROPE to incoming messages and find that this reduces the drop to 6.38 points on word labeling and 0.76 points on word grouping. Similar trend can be observed on FUNSD as well. Surprisingly, ROPE reduces performance drop more effectively than EdgeGeo on the larger payment dataset. Given these ablations, we conclude that reading order information is at least the same or more important than geometric features, and they bring orthogonal improvements to the overall performance.

Reading order encoding function. In practice, each target word usually has less than 8 neighboring words given a constructed β -skeleton graph. Therefore, a natural approach to assigning relative reading orders is to simply use the ROPE encoded indexes. In Table 2 we observe that simple index encoding immediately improves GCN without ROPE by 6.32 points on word labeling and 1.59 points on word grouping using payment corpus. Next we explore the popular sinusoidal function (with 3 base frequencies) for reading order encoding. It improves GCN without ROPE by 4.85 points on word labeling and 0.72 points on word grouping. Interestingly, sine function provides on par performance but does not outperform index encoding. The reason might be because the β -skeleton graph does not generate an extremely large number of neighbors, so simple index encoding is sufficient.

Sensitivity to OCR reading order. We investigate the robustness of ROPE to the quality of the input reading order. We shuffle the reading order provided by the OCR engine with a varying percentage of words before feeding into ROPE. Figure 4 exhibits the performance. For both word labeling and word grouping tasks, ROPE provides performance improvement up to less than 30% word or-

der shuffling on the large payment corpus. With 30% or more word order shuffled, we observe less performance degradation on the word labeling, suggesting that the word grouping task is more sensitive to the original OCR reading order.

5 Conclusion

We present a simple and intuitive reading order encoding method ROPE that is equivariant to relative reading order shifting. It embeds the effective positional encoding from sequence models while leveraging the existing spatial layout modeling capability of graphs. We foresee the proposed ROPE can be immediately applicable to other document understanding tasks.

Acknowledgements. We are grateful to Evan Huang, Lauro Beltrão Costa, Yang Xu, Sandeep Tata, and Navneet Potti for the helpful feedback on this work.

References

- Milan Aggarwal, Hires Gupta, Mausoom Sarkar, and Balaji Krishnamurthy. 2020. Form2seq: A framework for higher-order form structure extraction. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing (EMNLP)*.
- Brian Davis, Bryan Morse, Scott Cohen, Brian Price, and Chris Tensmeyer. 2019. Deep visual template-free form parsing. In *International Conference on Document Analysis and Recognition (ICDAR)*.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. BERT: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of the Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (NAACL-HLT)*.
- David Duvenaud, Dougal Maclaurin, Jorge Aguilera-Iparraguirre, Rafael Gómez-Bombarelli, Timothy Hirzel, Alán Aspuru-Guzik, and Ryan P Adams. 2015. Convolutional networks on graphs for learning molecular fingerprints. In *Advances in neural information processing systems (NIPS)*.
- Fabian B Fuchs, Daniel E Worrall, Volker Fischer, and Max Welling. 2020. Se (3)-transformers: 3d roto-translation equivariant attention networks. In *Advances in neural information processing systems (NeurIPS)*.
- Jonas Gehring, Michael Auli, David Grangier, Denis Yarats, and Yann N Dauphin. 2017. Convolutional sequence to sequence learning. In *International Conference on Machine Learning (ICML)*.

- Andrew Howard, Mark Sandler, Grace Chu, Liang-Chieh Chen, Bo Chen, Mingxing Tan, Weijun Wang, Yukun Zhu, Ruoming Pang, Vijay Vasudevan, et al. 2019. Searching for mobilenetv3. In *Proceedings of the International Conference on Computer Vision (CVPR)*.
- Guillaume Jaume, Hazim Kemal Ekenel, and Jean-Philippe Thiran. 2019. Funsd: A dataset for form understanding in noisy scanned documents. In *ICDAR-OST*.
- Anoop Raveendra Katti, Christian Reisswig, Cordula Guder, Sebastian Brarda, Steffen Bickel, Johannes Höhne, and Jean Baptiste Faddoul. 2018. Chargrid: Towards understanding 2d documents. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing (EMNLP)*.
- David G Kirkpatrick and John D Radke. 1985. A framework for computational morphology. In *Machine Intelligence and Pattern Recognition*.
- Frank Lebourgeois, Zbigniew Bublinski, and Hubert Emptoz. 1992. A fast and efficient method for extracting text paragraphs and graphics from unconstrained documents. In *International Conference on Pattern Recognition (ICPR)*.
- Xiaojing Liu, Feiyu Gao, Qiong Zhang, and Huasha Zhao. 2019. Graph convolution for multimodal information extraction from visually rich documents. In *Proceedings of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (NAACL-HLT)*.
- Xuanqing Liu, Hsiang-Fu Yu, Inderjit Dhillon, and Cho-Jui Hsieh. 2020. Learning to encode position for transformer with continuous dynamical model. In *International Conference on Machine Learning (ICML)*.
- Bodhisattwa Prasad Majumder, Navneet Potti, Sandeep Tata, James Bradley Wendt, Qi Zhao, and Marc Najork. 2020. Representation learning for information extraction from form-like documents. In *proceedings of the annual meeting of the Association for Computational Linguistics (ACL)*.
- Rasmus Berg Palm, Ole Winther, and Florian Laws. 2017. Cloudscan-a configuration-free invoice analysis system using recurrent neural networks. In *International Conference on Document Analysis and Recognition (ICDAR)*.
- Niki Parmar, Ashish Vaswani, Jakob Uszkoreit, Lukasz Kaiser, Noam Shazeer, Alexander Ku, and Dustin Tran. 2018. Image transformer. In *International Conference on Machine Learning (ICML)*.
- Yujie Qian, Enrico Santus, Zhijing Jin, Jiang Guo, and Regina Barzilay. 2019. GraphIE: A graph-based framework for information extraction. In *Proceedings of the Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (NAACL-HLT)*.
- Peter Shaw, Jakob Uszkoreit, and Ashish Vaswani. 2018. Self-attention with relative position representations. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (NAACL-HLT)*.
- Ilya Sutskever, Oriol Vinyals, and Quoc V Le. 2014. Sequence to sequence learning with neural networks. In *Advances in neural information processing systems (NIPS)*.
- Sandeep Tata, Navneet Potti, James B. Wendt, Lauro Beltrão Costa, Mark Najork, and Beliz Gunel. 2021. Glean: Structured extractions from templatic documents. In *International Conference on Very Large Data Bases (VLDB)*.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. In *Advances in Neural Information Processing Systems (NIPS)*.
- Petar Veličković, Guillem Cucurull, Arantxa Casanova, Adriana Romero, Pietro Lio, and Yoshua Bengio. 2018. Graph attention networks. In *International Conference on Learning Representations (ICLR)*.
- Renshen Wang, Yasuhisa Fujii, and Ashok C. Popat. 2021. Post-ocr paragraph recognition by graph convolutional networks. *arXiv preprint arXiv:2101.12741*.
- Liang Yao, Chengsheng Mao, and Yuan Luo. 2019. Graph convolutional networks for text classification. In *Proceedings of the AAAI Conference on Artificial Intelligence (AAAI)*.
- Michihiro Yasunaga, Rui Zhang, Kshitijh Meelu, Ayush Pareek, Krishnan Srinivasan, and Dragomir Radev. 2017. Graph-based neural multi-document summarization. In *Proceedings of the SIGNLL Conference on Computational Natural Language Learning (CoNLL)*.
- Jiaxuan You, Rex Ying, and Jure Leskovec. 2019. Position-aware graph neural networks. In *International Conference on Machine Learning (ICML)*.
- Wenwen Yu, Ning Lu, Xianbiao Qi, Ping Gong, and Rong Xiao. 2020. Pick: processing key information extraction from documents using improved graph learning-convolutional networks. In *International Conference on Pattern Recognition (ICPR)*.
- Chulhee Yun, Srinadh Bhojanapalli, Ankit Singh Rawat, Sashank Reddi, and Sanjiv Kumar. 2020. Are transformers universal approximators of sequence-to-sequence functions? In *International Conference on Learning Representations (ICLR)*.
- Xiaohui Zhao, Endi Niu, Zhuo Wu, and Xiaoguang Wang. 2019. Cutie: Learning to understand documents with convolutional universal text information extractor. In *International Conference on Document Analysis and Recognition (ICDAR)*.

01/14/99 15:06 FAX 200 023 0591 HAGENS BERMAN WA Fcc Reg. Aggr. 00017031

HAGENS BERMAN
130 FIFTH AVENUE, SUITE 1600, SEATTLE, WA 98101
TELEPHONE (206) 463-1792 FACSIMILE (206) 463-2524

FACSIMILE COVER SHEET

Date: January 14, 1999 No. of Pages: 37 (including this page)
From: Steve W. Berman File No.: 1129.01
Re: Tobacco Fee Payment Agreement and Release

COMMENTS:

Recipient(s)	Company	Phone No.	Fax No.
Mr. Meyer G. Koplow	Wachtel, Lipton, Rosen & Katz	(212) 403-1000	(212) 403-2000
Mr. Arthur F. Golden	Davis, Polk & Wardwell	(212) 450-4000	(212) 450-4800
Mr. Martin Barrington	Philip Morris Inc.	(917) 663-3199	
Mr. F. Anthony Burke	Brown & Williamson Tobacco Corp.	(502) 568-7297	
Mr. Ronald Milstein	Lorillard Tobacco Co.	(336) 335-7707	
Mr. Charles A. Blist	R.J. Reynolds Tobacco Co.	(336) 741-2998	
Mr. Stephen R. Patton	Kirkland & Ellis	(312) 861-2000	(312) 861-2200

Urgent! Deliver Immediately.

Please call the Support Center at (206) 268-9312 or not receive all of these pages or if there is a problem.

Professional Service Corporation

03573282

01/14/99 15:06 FAX 200 023 0591 HAGENS BERMAN WA Fcc Reg. Aggr. 00017031

HAGENS BERMAN
130 FIFTH AVENUE, SUITE 1600, SEATTLE, WA 98101
TELEPHONE (206) 463-1792 FACSIMILE (206) 463-2524

FACSIMILE COVER SHEET

Date: January 14, 1999 No. of Pages: 37 (including this page)
From: Steve W. Berman File No.: 1129.01
Re: Tobacco Fee Payment Agreement and Release

COMMENTS:

Recipient(s)	Company	Phone No.	Fax No.
Mr. Meyer G. Koplow	Wachtel, Lipton, Rosen & Katz	(212) 403-1000	(212) 403-2000
Mr. Arthur F. Golden	Davis, Polk & Wardwell	(212) 450-4000	(212) 450-4800
Mr. Martin Barrington	Philip Morris Inc.	(917) 663-3199	
Mr. F. Anthony Burke	Brown & Williamson Tobacco Corp.	(502) 568-7297	
Mr. Ronald Milstein	Lorillard Tobacco Co.	(336) 335-7707	
Mr. Charles A. Blist	R.J. Reynolds Tobacco Co.	(336) 741-2998	
Mr. Stephen R. Patton	Kirkland & Ellis	(312) 861-2000	(312) 861-2200

Urgent! Deliver Immediately.

Please call the Support Center at (206) 268-9312 or not receive all of these pages or if there is a problem.

Professional Service Corporation

03573282

11/12/99 15:02 FAX 000 000 0000 DAVIS POLK & WARDWELL

Fax Transmittal

480 Lexington Avenue
New York, NY 10017
212-450-4000

Sender: Charles Duggan
Date: November 12, 1999
Number of Pages (this page included): 6
Sender Voice Number: 212-450-4785
Sender Fax Number: 212-450-3785
Reference: 17556-002

If problems receiving this fax, call 212-450-4785

To	Fax Number	Company	Recipient Phone Number
Thomas M. Sobol	617-439-3278	Brown Rudnick Freed & Gesmer	617-330-8000
Joseph F. Rice	843-720-9290	Ness, Motley, Loadholt, Richardson & Poole	843-720-9000
Robert V. Costello Jeffrey D. Woolf	617-722-0286	Schneider, Reilly, Zabin & Costello	617-227-7500
Richard M. Helmann	415-956-1008	Lieff, Cabraser & Helmann	415-956-1000
Michael P. Thornton	617-720-2445	Thornton, Early & Naumes	617-720-1333

Message:

03573282

11/12/99 15:02 FAX 000 000 0000 DAVIS POLK & WARDWELL

Fax Transmittal

480 Lexington Avenue
New York, NY 10017
212-450-4000

Sender: Charles Duggan
Date: November 12, 1999
Number of Pages (this page included): 6
Sender Voice Number: 212-450-4785
Sender Fax Number: 212-450-3785
Reference: 17556-002

If problems receiving this fax, call 212-450-4785

To	Fax Number	Company	Recipient Phone Number
Thomas M. Sobol	617-439-3278	Brown Rudnick Freed & Gesmer	617-330-8000
Joseph F. Rice	843-720-9290	Ness, Motley, Loadholt, Richardson & Poole	843-720-9000
Robert V. Costello Jeffrey D. Woolf	617-722-0286	Schneider, Reilly, Zabin & Costello	617-227-7500
Richard M. Helmann	415-956-1008	Lieff, Cabraser & Helmann	415-956-1000
Michael P. Thornton	617-720-2445	Thornton, Early & Naumes	617-720-1333

Message:

03573282

Figure 6: Sample output of the word grouping task on FUNSD with a few failure cases.