Deep Transfer Learning & Beyond: Transformer Language Models in Information Systems Research

ROSS GRUETZEMACHER, Wichita State University, W. Frank Barton School of Business

DAVID PARADICE, Auburn University, Harbert College of Business

AI is widely thought to be poised to transform business, yet current perceptions of the scope of this transformation may be myopic. Recent progress in natural language processing involving **transformer language models (TLMs)** offers a potential avenue for AI-driven business and societal transformation that is beyond the scope of what most currently foresee. We review this recent progress as well as recent literature utilizing text mining in top IS journals to develop an outline for how future IS research can benefit from these new techniques. Our review of existing IS literature reveals that suboptimal text mining techniques are prevalent and that the more advanced TLMs could be applied to enhance and increase IS research involving text data, and to enable new IS research topics, thus creating more value for the research community. This is possible because these techniques make it easier to develop very powerful custom systems and their performance is superior to existing methods for a wide range of tasks and applications. Further, multilingual language models make possible higher quality text analytics for research in multiple languages. We also identify new avenues for IS research, like language user interfaces, that may offer even greater potential for future IS research.

CCS Concepts: • **Information systems** → **Data mining**; **Information retrieval**;

Additional Key Words and Phrases: Natural language processing, text mining, artificial intelligence, deep learning, transfer learning, language models

ACM Reference format:

Ross Gruetzmacher and David Paradice. 2022. Deep Transfer Learning & Beyond: Transformer Language Models in Information Systems Research. *ACM Comput. Surv.* 54, 10s, Article 204 (September 2022), 35 pages. <https://doi.org/10.1145/3505245>

1 INTRODUCTION

There is tremendous hype about **artificial intelligence (AI)** and its potential to transform business. However, many organizations have struggled to see real benefits to their bottom lines due to AI initiatives [Fountain et al. 2019]. While Fountain et al. are correct to suggest that organizations need to change their culture to reap the benefits of AI, it is also true that many of the benefits of AI have yet to be realized because the technology is still in its nascency and research progress continues at a rapid pace. There is no apparent reason to suspect this progress to slow, either, and leading

Authors' addresses: R. Gruetzmacher, Wichita State University, W. Frank Barton School of Business 1845 Fairmount St. Wichita, Kansas, 67260 USA; email: ross.gruetzmacher@wichita.edu; D. Paradice, Auburn University, Harbert College of Business 405 W. Magnolia Ave Auburn, Alabama, 36849 USA; email: dparadice@auburn.edu.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

© 2022 Copyright held by the owner/author(s). Publication rights licensed to ACM.

0360-0300/2022/09-ART204 \$15.00

<https://doi.org/10.1145/3505245>

organizations in business consulting, economics and policy all foresee AI-driven transformative change in business on the horizon.

Rapid progress in the use of deep learning – the AI technique driving current progress – for image processing and speech recognition in the early-to-mid-2010s was impressive, and progress in deep reinforcement learning has drawn a lot of media attention by demonstrating superhuman performance in a number of games [Silver 2017; LeCun 2015]. Yet, it is debatable as to whether the perceived progress is living up to the hype in practice. While deep learning certainly has valuable applications in business operations and business analytics [Kraus 2020], it has not yet led to significant productivity gains [Brynjolfsson 2020].

However, there is reason to think that recent progress in the use of pretrained language models, which emerged in 2018, may be different. Anya Belz, in the opening keynote of the 2019 International Conference on Natural Language Generation, only days after the release of the most powerful generative language model to date, T5 [Raffel 2020], openly asked “Did T5 just solve general natural language generation?” [Belz 2019]. This question was not made in jest, rather, it was delivered with a sense of dismay; progress truly is being made at a rate which many who have worked on the problem for a long time find unnerving. T5 is no longer the most powerful generative language model, and its successor, GPT-3 [Brown 2020], may have improved even further beyond T5 than it had improved beyond its predecessors.

The three primary reasons to believe that this recent progress is different are not evidenced by the nature of the progress alone but in large part by the nature of how these techniques naturally fit into organizations’ operations. *First*, organizations create and collect large amounts of unstructured data. This data is widely thought to contain information that, if harnessed, could be very valuable. For this reason, even moderately effective text mining techniques are already able to deliver tremendous value to organizations in numerous domains ranging from policy [Ngai 2016] to finance [Kraus 2017] or biomedical engineering [Gonzalez 2015]. *Second*, this new generation of pretrained language models harnesses the enormous potential of unsupervised and semi-supervised learning [Collobert 2008]. This means that these models can be initially trained in an unsupervised fashion on very large corpora and then later they can be fine-tuned (i.e., **deep transfer learning**) on an organization’s labeled, task-specific data so that they outperform existing text mining techniques for a variety of tasks specific to the organization’s needs. *Third*, progress in this area is showing no signs of slowing down, and more advanced capabilities from increasingly powerful systems may continue for some time [Kaplan 2020]. Examples include chatbots’ capabilities which are likely to bring long anticipated **language user interfaces** (LUIs) [Brennan 1991] to a wide variety of human-computer interactions, and **few-shot learning** capabilities that can reduce the training and skills required for using language models while creating the possibility of truly novel applications.

This paper is intended for researchers and practitioners who are interested in any type of business research that may benefit from analysis of large amounts of unstructured text data (e.g., emails, reviews, social media posts) as well as those interested in applications of LUIs, both practically and in information systems (IS) research. This study makes several contributions:

- (1) It identifies and reviews the state-of-the-art literature for a powerful application of deep learning that has not yet been effectively incorporated into the toolbox of IS researchers.
- (2) It conducts a literature review of existing work in leading IS journals using text mining, clearly identifying limitations of existing work and the benefits of using the new tools.
- (3) It proposes concrete research ideas that go beyond simply improving existing work to offer new directions for researchers and practitioners to explore.

No text mining or NLP experience is necessary for readers of this article¹, but we do assume that readers are familiar with the concepts of neural networks and deep learning. In the remainder of the paper we first survey the recent progress that has led to these powerful new tools. We next survey extant applications of text mining for business analytics and IS research. We then consider recent applications of these new tools within a discussion of their implications for both research and practice. We follow the discussion by summarizing its salient elements, including the most promising avenues for future work, and finally we leave concluding remarks.

2 RECENT PROGRESS IN NEURAL LANGUAGE MODELS

The subfield of machine learning known as computational linguistics or **natural language processing (NLP)** has been one of the primary focuses for AI researchers since the beginning of the study of AI: the first conference on machine translation preceded even the 1956 Dartmouth workshop, thought of as a seminal event of the field, and the necessity of NLP for AI was clear as early as Turing's proposed test for intelligence (a.k.a. the Turing test) [Turing 1950]. The early years of NLP research (i.e., 1960-1985) centered on what is known as the rationalist approach. Statistical NLP, which takes an empiricist's approach, did not become the dominant school of thought until the 1990s [Manning 1999]. Statistical NLP assumes that a large degree of latent semantic knowledge resides in text corpora, and, in order to encode this knowledge, numerical representations of language are necessary [Smith 2020]. Such numerical representations of words are called **word representations** (a.k.a. word vectors or word embeddings), and they are a fundamental building block of statistical NLP.

Originally, words were encoded simply by assigning an integer to each unique token. However, integers are poor word representations because they do not allow semantic information to be shared across words with similar properties [Smith 2020]. Distributed representations, on the other hand, can contain continuous values for each dimension, and these dimensions can be thought of as semantic features of the word being represented capable of encoding the semantic relations among words [Senel 2018]. For example, if we assign a dimension of a word vector to be associated with weight (in grams), *feather* might have a value of 0.001, *penny* might have the value of 2.5 and *car* might have a value of 1,250,000.0, but adjectives like *green* and *chilly* would have values of zero. Creating representations of words in this way is known as feature engineering, but it is not practical for most corpora, and surely not for an entire language. In most cases, learning word representations is far more useful. Early semantic representations utilized frequency-based methods like singular value decomposition of the co-occurrence matrix. Such approaches power many widely used text mining techniques like **latent Dirichlet allocation (LDA)** [Blei 2003]. These techniques comprise one family of word representations called global matrix factorization models [Pennington 2014].

Using word representations to encode the semantics of words in a language, **statistical language models** (a.k.a. probabilistic language models or simply language models) can be created to model the probability of word occurrence in sentences. Specifically, language models are probability distributions of sequences of words that are useful for problems that require the prediction of the next word in a sequence given the previous words. n-gram models are a very simple form of language model that are commonly used in text mining, and such simple models have long been used for a variety of other tasks including spell check, machine translation and speech and handwriting recognition [Manning 1999].

¹We do not go into technical details here as our purpose is to inform readers about the possible applications of these techniques, both now and in the future, relative to existing text mining techniques.

2.1 Neural Word Representations

Neural language models comprised of distributed word representations were first proposed as a solution for the curse of dimensionality [Bengio 2003]. Collobert and Weston [2008] then demonstrated the value of using deep learning for learning distributed representations of words from large unlabeled corpora, then transferred the learnt knowledge to multiple tasks learned simultaneously through further training (i.e., fine-tuning) on labeled datasets (i.e., deep transfer learning). Early last decade, Collobert et al. [2011] described the first pretrained neural word representations that were able to achieve strong performance on major NLP tasks.

In the time since these early studies, distributed word representations have become widely preferred over alternate representations. One reason feature engineering is not practical for NLP is due to the challenge of identifying all of the relevant features of words that would need to be represented in order to capture the entire semantics of a corpus or language. However, representation learning generates a latent feature space where features are not constrained by the need to map directly to human concepts in natural language, which makes it much easier to capture the rich semantics of a language with a limited number of features (e.g., 100 to 300).

The first strong neural language model to learn practically useful word representations in this manner was word2vec [Mikolov 2013b]. For training, Mikolov et al. proposed two different architectures: one for predicting the current word based on context (i.e., **continuous bag-of-words or CBOW**) and another for predicting the surrounding words given the current word (i.e., continuous skip-gram). The former was better suited for small corpora while the latter was better suited for scaling to large corpora. The new techniques proposed by Mikolov et al. were able to generate rich word representations that captured fine-grained semantic and syntactic regularities better than previous models. This led to the widespread use of word representations in NLP.

The ability of these neural word representations to explicitly encode numerous linguistic regularities and patterns exhibited some very interesting characteristics: the relationships between two words could be represented as linear translations and that simple vector operations could be used to evaluate the concepts of semantic and syntactic similarity between words [Mikolov 2013a]. For example, the vector operation $\overrightarrow{Madrid} - \overrightarrow{Spain} + \overrightarrow{France}$ was closer to $\overrightarrow{Paris}$ than any other word. Even vector addition alone had valuable results: in the latent feature space $\overrightarrow{Germany} + \overrightarrow{capital}$ was close to $\overrightarrow{Berlin}$ and $\overrightarrow{Russia} + \overrightarrow{river}$ was close to $\overrightarrow{Volga River}$.

word2vec [Mikolov 2013b] was the first in a new family of word representation models that are very useful for analytics because it enabled building custom models from rich, learnt semantic features. However, it was not without limitations. For example, it did not leverage document level information during training. Alternately, global matrix factorization models were able to leverage document level statistical information but were unable to perform well on the analogical evaluation in which word2vec excelled (i.e., vectors' semantic and syntactic similarity). Pennington et al. [2014] attempted to address these issues with GloVe (global vectors for word representation), which was able to make efficient use of document level statistics like global matrix factorization methods while also generating representations with a meaningful vector space substructure like that of word2vec. GloVe outperforms word2vec on some benchmarks, but both techniques are still widely used.

We often refer to word representations like word2vec and GloVe [Mikolov 2013b; Pennington 2014] as pretrained word representations because these word representations are openly available for public download having already been pretrained on large text corpora. When used in this fashion they can be very powerful because the Web-based corpora that they have been trained on

are often too large for training by independent researchers with limited computational resources. Even for those with the requisite computing power, these pretrained representations save time from tuning hyperparameters and cleaning corpora. However, it is also common practice to train representations on smaller, domain-specific corpora which require fewer computational resources. This can be worth the time spent tuning hyperparameters due to the improvements that can be obtained by using domain-specific word representations.

2.2 Neural Language Models

Traditional deep neural networks² are not well suited for language processing because NLP tasks often require mapping from vectors of different lengths. For example, a language model designed to predict the next word in a sentence must operate on an input of one word as well as an input of 20 words in order to predict the next word. **Recurrent neural networks (RNNs)** are neural networks with feedback connections that are suitable for machine learning tasks requiring such sequence-to-sequence mapping [Murphy 2012]. Their suitability for these tasks is due to the fact that, unlike normal neural networks, they are able to map from vectors of varying lengths to other vectors of varying lengths. One of the challenges that RNNs face in NLP is known as the vanishing gradient, which can cause training to fail, but there are techniques that can be used to mitigate this problem. **Long short-term memory (LSTM)** [Hochreiter 1997] models are a form of RNN that use a gating mechanism to address this problem, and they have long been used for a large number of NLP applications. An LSTM gating mechanism controls the flow of information to hidden neurons, which, for sequence-dependent data (e.g., sentences), can encode the meaning of a sequence while remembering (or forgetting) the most (or least) salient elements. Other approaches can also be applied to RNNs to counteract the vanishing gradient problem [Mikolov 2014], such as the gated recurrent unit [Chung 2014], but none have been as effective as the LSTM. LSTMs work very well for a variety of challenging tasks; however, LSTMs typically rely on supervised learning and require a unique labeled training dataset for each task. They are well-suited for a variety of NLP tasks ranging from classification to translation to text generation and were the dominant NLP technique prior to the development of attention and the transformer architecture.³

All of the word representations described thus far are static word representations, and they all have one major limitation: they attempt to represent words in all possible contexts with a single vector. However, words have different meanings in different contexts and thus are not always best represented in a static manner. **Contextual word representations (CWRs)** offer a solution for this based on the premise that if each word is going to have a unique representation then each vector should be dependent on a separate context vector representing the sequence of nearby words. CWRs were popularized by Peters et al. [2018] with the ELMo language model. ELMo was significant in that it demonstrated **state-of-the-art (SOTA)** performance on not just one NLP task but six, suggesting that performance gains from this approach were likely for a wide variety of NLP tasks. It also ushered in a class of pretrained language models with rich word representations embedded in the weights. *Because these new language models are both language models and (able to generate) rich word representations, we refer to them simply as language models.*

²Deep learning is a form of representation learning: a type of machine learning involving learning of representations or features in data. Mathematically, it can be thought of as a technique for learning a function to map from the input data to the output data.

³While LSTMs are thought to have been the dominant technique prior to the transformer, convolutional neural networks were (and are) still very capable and even preferable for many applications (e.g., classification). Recent work [Tay 2021] suggests that convolutional neural networks may still have significant capabilities despite the recent prevalence of transformer-based language models.

2.3 Deep Transfer Learning

Transfer learning has long been a topic of interest for machine learning researchers. In fact, it predates deep learning and is a machine learning problem unto itself. However, deep transfer learning is very powerful and has become a topic of interest, being used for tasks from image processing to NLP.

Generally speaking, transfer learning involves the transfer of knowledge learned from one learning task to improve results or speed up training for another task. It effectively removes the need to train a model from scratch by enabling specialized training for a new task via the fine-tuning of an existing, pretrained model. **Deep transfer learning (DTL)** refers to this use of deep learning for pretraining models on large amounts of data, either from labeled data for supervised pretraining or from unlabeled data for unsupervised pretraining. These pretrained models can then be fine-tuned on task-specific datasets to transfer the knowledge learnt from the more general, original training datasets to the domain-specific applications and use cases.

Pretraining, as it is most commonly used for learning word representations, is a specific form of unsupervised learning known as self-supervised learning. Self-supervised learning does not require explicit labels for data as supervised learning techniques do, rather, implicit supervisory signals from the data are autonomously extracted and used during pretraining. In the case of NLP, these signals come from the sequence of the words, e.g., a model can be trained by masking a single word and training the model to predict that word given the surrounding words. Pretraining is critical to DTL model performance, and there are several key aspects that are important to understand. Typically, pretraining is performed using very large corpora, and corpora selection or curation can have a significant impact on model performance and end tasks. Also, the selection of a pretraining objective and the approach for self-supervision can have significant impacts on model performance and end tasks. The relevance of these elements will be explained further in the following sections.

Fine-tuning is a critical component of DTL, too. It refers specifically to the further training of a pretrained model on a smaller, labeled dataset. For NLP it refers specifically to the process of leveraging the vast semantic knowledge contained in large, pretrained models for application to domain-specific tasks involving small domain-specific datasets. It is valuable because it enables the simple development of custom, SOTA NLP systems for a great variety of tasks with relatively small, labeled datasets and with significantly less effort than previous techniques (e.g., LSTMs). Oftentimes with the latest language models SOTA performance for a task can be attained by simply fine-tuning on task-specific datasets [Howard 2018].

In the context of language models, transfer learning is considered to have four steps: pretraining, further pretraining, pre-finetuning and fine-tuning. The first and last steps have been explained, but the steps in-between can be useful for significantly improving performance when using DTL. Further pretraining involves pretraining an already pretrained model on an alternate dataset, commonly one which is smaller and either domain-specific or task-specific. While more computationally expensive than fine-tuning, further pretraining is still much less resource intensive as pretraining the model from scratch and can be worth the cost when end task performance is critical. Additionally, further pretraining differs from initial pretraining in that further pretraining does not involve self-supervised signals and is not full self-supervised learning like initial pretraining.

Like further pretraining, pre-finetuning is performed on an already pretrained model in order to further refine representations prior to end-task fine-tuning. It involves the use of a broad supervised dataset for multitask training in order to encourage learning representations that will generalize better to a variety of downstream tasks [Aghajanyan 2021]. Less computationally expensive than further pretraining, pre-finetuning can be used to improve performance when end-task effectiveness is critical or to improve zero-shot performance [Wei 2021].

2.4 Transformer Language Models

Because the study focuses on transformers and transformer-based language models, we split this section into two subsections: a high-level overview of recent progress, consistent with the overall narrative, and a discussion of more technical details about the transformer and models based on it that are relevant to IS researchers.

2.4.1 Overview. In late 2017 the transformer architecture was first proposed by Vaswani et al. [2017]. At the time LSTMs were the prevailing paradigm in NLP, but they did not work terribly well or reliably for very long sequences or transfer learning. The transformer presented a novel way to incorporate an attention mechanism in deep feedforward networks that allowed it to capture long range sequence dependencies like the LSTM, but with a larger context window for longer sequences. The transformer was also easily parallelizable and highly scalable. Due to this, transformers have been trained using unprecedentedly large corpora [Raffel 2020; Brown 2020]. We distinguish language models using the transformer as **transformer language models (TLMs)** because they perform remarkably better than LSTM-based models and they scale very well [Kaplan 2020].

There are a large number of TLMs, but here we will initially focus on three of the most significant with respect to practicality, novelty and improvement upon previous models.⁴ TLMs first emerged over summer of 2018 [Radford 2018], but the most powerful early model was **BERT (bidirectional encoder representations from transformers)**⁵ [Devlin 2018], which was demonstrated in late 2018. In the time since it has become the most widely used TLM⁶, and is able to achieve SOTA performance on a wide number of tasks due to its versatility.

The next major model advance was the **text-to-text transfer transformer (T5)** [Raffel 2020], which was developed specifically for transfer learning and is designed to operate solely through text generation by framing all text-based language problems as text-to-text tasks.⁷ We refer to language models like this that operate solely through text generation as generative language models. In contrast to models like BERT where fine-tuning involves adding a fully connected layer and output neurons, which means that separate models are necessary for multiple tasks, T5 is intended to be fine-tuned on multiple tasks by default. By design T5 can be trained on multiple tasks simultaneously, in the fashion proposed by Collobert and Weston [2008] over a decade earlier.

The final language model we mention here is also a generative language model: the **generative pretrained transformer 3 (GPT-3)** [Brown 2020]. GPT-3 is an OpenAI TLM which uses the same architecture as its predecessor, GPT-2 [Radford 2019]. What makes GPT-3 unique is its scale: GPT-3 was scaled to a model size, measured by number of parameters, an order of magnitude larger than any previous model and was pretrained on the largest dataset to date. This required

⁴TLMs are trained in a variety of model sizes (number of parameters). For this survey we only consider the largest of each TLM.

⁵BERT has two notable characteristics: it is trained bidirectionally and it is a masked language model. Instead of being trained for predicting the next word in a sentence it is trained to predict missing words in a sentence. It does so by masking (i.e., masked language model) 15% of the words and training bidirectionally to predict the missing words as well as the next sentence.

⁶There are now a large number of variants of BERT, either architectural variations or models that were pretrained on a domain-specific dataset. BERT is so popular that a survey [Xia 2020] on the different variants and how to pick the best one for different types of problems was published recently at one of the premier NLP conferences. (This survey is a valuable resource).

⁷T5 was the result of a large-scale study by Google on the limits of transfer learning from transformer language models and it is not like previous models because it operates as a text-to-text language model, meaning that it both receives text as an input and produces text as an output. Most NLP tasks can be formulated in this manner, and this enables T5 to train a single model to perform multiple tasks during inference by appending a label associated with each unique task to the beginning of the input text.

extreme investments in computational resources and distributed computing infrastructure, but led to surprising improvements on zero-, one- and few-shot learning tasks.

Few-shot learning refers to the ability of a system to be able to learn without the need for even modestly sized datasets typically used for fine-tuning. For example, being a generative language model, the model could be trained by providing (k) questions as input and the (k) correct answers as the training targets (e.g., for one-shot learning k is one). For two-digit addition problems GPT-3 achieves 99.6% accuracy with only one example, and with no examples – i.e., zero-shot learning – GPT-3 still achieves 76.9% accuracy. Few-shot learning would involve k training prompts and targets, with a limit set by the fixed context window of 2,048 tokens. Thus, if the model does not perform well on tasks with zero- or one-shot learning, more examples can be used to improve performance. Further work from OpenAI suggests language model performance will continue to scale with computational resources and dataset size, with no plateau in sight [Kaplan 2020].

2.4.2 Under the Hood. The original transformer proposed by Vaswani et al. [2017] includes both encoder and decoder components. The encoder encodes the input sequence into a high dimensional feature space, and the decoder converts high dimensional representations back into words. This is known as a **sequence-to-sequence (seq-to-seq)** model, and, as such, it is naturally well suited for tasks like machine translation. However, the key contribution of the paper was not the encoder-decoder element, but rather that, unlike the LSTM, transformers do not use recurrence or require any sequential computation. Thus, transformers are not subject to the vanishing gradient problem.⁸ Attention mechanisms predate the transformer [Xu 2016], but the transformer made practical use of attention in a novel and powerful way that enabled SOTA results on a major machine translation benchmark.

Pretraining is the most critical element of a TLM as this is what distinguishes different TLMs. However, TLMs typically fall into one of three categories with respect to their pretraining: autoencoding, autoregressive or seq-to-seq [Wolf 2019]. BERT is an example of an autoencoding TLM whereas GPT-3 is an example of an autoregressive TLM, and the original transformer is an example of a seq-to-seq TLM. The distinction between models is determined by the pretraining scheme. For example, BERT encodes documents bidirectionally and replaces random tokens with masks, then is trained to predict masked tokens as well as the next sentence. This contrasts with GPT-3, where tokens are predicted autoregressively with a left-to-right decoder. Autoencoding models perform best at discriminative tasks (e.g., classification, regression; tasks where BERT excels) while autoregressive models perform best at generative tasks (e.g., summarization, dialogue; where GPT-3 excels).

BERT is an example of a seq-to-seq model that attempts to bridge the divide between autoencoding and autoregressive models [Lewis 2020a]. It used a variety of denoising schemes for pretraining its denoising autoencoder. The results suggested that new pretraining schemes could lead to strong performance on generative tasks without sacrificing performance on discriminative tasks, and that different approaches to corrupting documents⁹ during pretraining may be better suited for specific downstream tasks. Another seq-to-seq TLM based on a new type of denoising autoencoder, MARGE [Lewis 2020b], utilizes an alternative self-supervision technique to the dominant token masking paradigm; similar documents from other languages are used to assist in reconstruction

⁸While the LSTM uses a gating mechanism to mitigate the problem of the vanishing gradient, this does not overcome it completely, and the vanishing gradient still limits the context window of the LSTM. By not using recurrence at all, the transformer avoids this problem which results in a larger context window and the transformer's most significant improvement over the LSTM.

⁹By corrupting we are referring to pretraining approaches such as masking 15% of input tokens during training, as with BERT. A variety of corruption approaches are explored with BERT, and different approaches could have impacts on downstream tasks with fine-tuning.

of the input document. MARGE performs strongly on a wider range of tasks in many languages – both discriminative and generative – than previous models.

2.5 Natural Language Understanding

Natural language understanding (NLU) is typically thought to be a more general, longer-term goal for NLP researchers. We mention it here because significant effort has been made to develop measures to quantify progress in this domain, and these measures demonstrate the recent progress of TLMs. In April of 2018, researchers from leading institutions in business and academia realized the need for a new means of assessing progress and developed the **General Language Understanding Evaluation (GLUE)** benchmark [Wang 2018], which was intended to be a benchmark for measuring progress toward NLU. Just over a year later, in June of 2018, Microsoft had surpassed the human baseline for GLUE [Liu 2019a]. However, this was anticipated, and a more difficult SuperGLUE benchmark was released [Wang 2019a].

BERT was used as the initial baseline for the SuperGLUE benchmark achieving a score of 69.0, well below the human baseline of 89.8, but less than three months later a team from Facebook AI Research had demonstrated a robustly optimized version of BERT (RoBERTa¹⁰) that was able to achieve a SuperGLUE score of 84.6¹¹ [Liu 2019c]. This striking progress led to speculation that, like progress in other domains such as self-driving vehicles, the first 95% of the task was less difficult than previously perceived, but that the last 5% would become exponentially more challenging. However, less than three months later, and to the dismay in the natural language generation community [Belz 2019], T5 was released demonstrating a score of 88.9¹² on SuperGLUE – within a point of the human performance baseline [Raffel 2020].

Following the release of T5 [Raffel 2020], progress appeared to slow for six months, which seemed to suggest that early intuition about the last 5% becoming more difficult may be valid. However, this lag was again shown to be unfounded by GPT-3 [Brown 2020]. *GPT-3 impressed for many reasons, but its performance might be best summarized by considering that it achieved a SuperGLUE score of 71.8, over 4% higher than BERT, simply by using few-shot learning ($k = 32$).* As impressive as this is, it is also important to consider that the variance among scores for the different tasks comprising the aggregate measure was dramatically higher for GPT-3. For example, BERT outperformed GPT-3 by over 45% on one complex linguistic task but GPT-3 outperformed BERT by over 30% on a widely used causal reasoning benchmark.

T5, with the use of an **Unsupervised Data Generation (UDG)** procedure [Wang 2021b], has since surpassed human-level performance on SuperGLUE. Other models have also exceeded human-level performance on SuperGLUE, including a decoding-enhanced BERT with disentangled attention (DeBERTa) from Microsoft [He 2021], which builds on RoBERTa with disentangled attention and enhanced mask decoder training. Another multilingual model, ERNIE 3.0 [Sun 2021], from Baidu, has even surpassed the SuperGLUE performance of T5. To achieve this feat, ERNIE 3.0 fuses an auto-regressive network and an auto-encoding network for training on both text data and a large-scale knowledge graph.

One research topic that clearly falls under the purview of NLU is chatbots. Progress in this area has also seen great improvements since 2018, with the most significant progress occurring recently. At the beginning of 2020 Adiwardana et al. [2020] demonstrated Meena, an open domain chatbot that was thought to have more “human-like” conversation as measured by a novel metric

¹⁰RoBERTa was simply a replication study of BERT that explored the significance of different hyperparameter choices and training dataset size. They found that doing away with the next sentence prediction in pretraining, and some other training modifications, greatly improved the performance of BERT.

¹¹RoBERTa should be used instead of BERT when possible due to the easy performance gains from pretraining.

¹²In the time since, T5 has been retrained to score even higher on SuperGLUE with an 89.4.

that required evaluation of responses through crowdsourcing. Only a few months later, Facebook AI released Blender Bot open source [Roller 2020a], which was significant as Blender Bot outperformed Meena on human judged evaluations. Because Blender Bot is a TLM-based chatbot, it can be fine-tuned for domain-specific tasks, offering new opportunities for IS researchers.

2.6 Ongoing Research and Future Directions

The review thus far brings readers up-to-speed with respect to where the SOTA is for TLMs and NLP. This section discusses some focal areas of current research where substantial progress is being made that stands to dramatically improve future capabilities beyond what might be anticipated from the research discussed prior.

One of the most pressing limitations of TLMs is that they are terribly inefficient; the largest models can only be fine-tuned on the cloud, which can be costly for researchers, and such models are not possible to use on edge or mobile devices. However, important work has been conducted to address this. For example, DistilBERT [Sanh 2019] is 60% faster and 40% smaller than BERT but still retains 74% of BERT's performance above the GLUE baseline. Other models focus on efficiency for specific tasks, such as TopicBERT, which is intended for document classification and achieves a 40% speedup while retaining 99.9% of BERT's performance on five tasks [Chaudhary 2020]. Further, new work from Schick and Schütze [2020] demonstrates strong few-shot learning performance on SuperGLUE using ALBERT (a lite BERT) [Lan 2019], one of the most efficient TLMs.

When Vaswani et al. [2017] first demonstrated the transformer, it had SOTA results on a major machine translation benchmark. TLMs work well for translation because they can represent the semantics of multiple languages in a shared, high-dimensional latent space. However, there are other applications of multilingual TLMs, and recent work has begun focusing on massively multilingual TLMs trained on eight languages or more, for which new multi-domain and multi-task benchmarks have been developed [Siddhant 2020; Liang 2020]. The most widely used multilingual TLM is XLM-R, which is pretrained using text from 100 languages [Conneau 2020]. Optimizing these models for multiple tasks across multiple languages is one of the biggest challenges for multilingual TLMs, but the latest work has demonstrated substantial progress in this direction [Wang 2021a].

While revolutionary, the transformer is far from perfect. One significant limitation is the fixed length context window, which, while far greater than the LSTM, could be even more useful if extended. Early work addressing this, the Transformer-XL [Dai 2019], is used in XLNet [Yang 2019b], which demonstrates superior performance on general sentiment analysis. Other recent work has demonstrated a TLM with a context window of up to one million words [Kitaev 2020]. Another major limitation of transformers is the poor scalability of the self-attention mechanism. To address this, Zaheer et al. [2020] have proposed BigBird, a sparse attention mechanism that drastically improves upon the original transformer's attention mechanism, which could have practical implications for tasks such as longer document summarization and question answering. Choromanski et al. [2021] have gone further proposing another sparse attention mechanism that demonstrates generalized attention which may lead to even greater improvements once used for training large TLMs.

While the transformer has been exploited extensively as the current architecture of choice for the pretrain/fine-tune paradigm, recent work has shown that convolutional neural networks show promise in this area as well [Tay 2021]. This work suggests that architectural progress should be differentiated from progress in pretraining, and that convolutional neural networks outperform TLMs in some cases. Thus, further work is merited to explore alternative architectures to the transformer¹³ within the pretrain/fine-tune paradigm.

¹³This survey focuses on TLMs because only TLMs alone have demonstrated tremendous progress in NLP – at an equivalent level to AlexNet [Krizhevsky 2012] in computer vision last decade – however, if alternative architectures are as successful, the directions for future IS research suggested in later sections would still apply.

Finally, while much progress has been made in the domain of NLU, it is clear that there are limitations to distributional semantics (i.e., purely probabilistic models of semantic similarity). Grounded semantics refers to the grounding of semantic concepts to knowledge learnt from other forms of data (e.g., images, video, simulation), and is thought to be the next step toward NLU. Early work in this direction has focused on improving “commonsense reasoning” through grounding CWRs via training a multimodal model on a question answering dataset about images from movie scenes [Zellers 2019]. Recent work from Tan and Bansal [2020] builds on their previous work on cross-modal TLMs (LXMERT) [Tan 2019] by exploring the possibility of a visually-supervised TLM through a process they call “vokenization.”

Visual grounding may produce impressive results at present, but effective communication relies on a shared understanding of the world, one which is learnt from experience. Drawing from this notion, Bisk et al. [2020] have proposed five levels of *World Scope*: (1) corpus (the past), (2) the Web (most of current NLP), (3) perception (multimodal NLP), (4) embodiment (situated action taking), and (5) the social world. This new perspective is exciting because it situates existing work and identifies the next steps forward toward the shared understanding of the world necessary for truly successful linguistic communication through language grounding.

2.7 Summary

After the preceding overview of statistical NLP up to the current state of the field, we want to highlight three distinct periods through which recent progress can be better understood. The first period began with the inception of the field and lasted until 2013. This period was characterized in large part by hand crafted features and to some degree the emergence of statistical NLP. The bulk of text mining techniques used in practice today originated in this period. The second period began in 2013 with the word2vec model [Mikolov 2013b]. This ushered in a new, data-driven period characterized by neural word representations and neural language models. The most recent period began in 2018 and involves the topics which are the focus of this study.

This current era of NLP is defined by three significant components: (1) CWRs, (2) DTL/fine-tuning, and (3) TLMs. Combined, these elements are transforming the study of NLP because they enable leveraging unsupervised learning for a large number of tasks and applications. Fundamentally, this means that there is no limit to the amount of data that can be used for training, which brings new meaning to the phrase “big data”. Further, gains from continued scaling of language models are not expected to plateau anytime soon [Kaplan 2020]. Such continued progress combined with grounded semantics advances may usher in truly transformative language processing, and it is essential for IS researchers to be familiar with progress in this domain.

3 TEXT MINING IN INFORMATION SYSTEMS RESEARCH

As this paper explores the future applications of TLMs in IS research, we felt it appropriate to conduct a thorough review of recent work in IS that used text mining or NLP. Due to the widespread application of these techniques, we limited our review to articles published (and preprints accepted for publication) from 2016 to 2020 in the three leading IS journals: Information Systems Research, the Journal of MIS and MIS Quarterly.^{14,15} We chose to focus our review on three terms:

¹⁴Even more articles using text mining appeared in other IS outlets such as Decision Support Systems and the International Conference on Information Systems proceedings. However, only articles from the top three IS journals were selected for inclusion in order to reduce noise because there were so many articles from these other outlets and the articles in the top three journals were deemed to be most representative of rigorous IS research. Moreover, very few articles using text mining appeared in the other basket of eight IS journals.

¹⁵It is possible that some articles slipped through, like Benjamin et al. 2016 who only mention *text classification* and none of our keywords.

“text mining,” “sentiment analysis” and “natural language processing.” We queried these terms using Google Scholar’s advanced search feature on September 10th, 2020. Only work that utilized these methods as part of a model were included in the results – studies that just mentioned the terms or only used them for robustness checks were discarded.

This process resulted in 55 papers meeting our criteria. These papers are listed in Table 1 which depicts relevant features of each study. Each paper was carefully reviewed and coded with respect to these relevant features, seven of which characterized the techniques and seven more that characterized other relevant aspects of the studies. Each author independently coded each paper, and for disagreement, a discussion was conducted to arrive at a consensus. For coding we classified each feature with check marks of two shades to represent weak correspondence (grey) and strong correspondence (black). The four results that cited the two original neural word vector models [Mikolov 2013b; Pennington 2014] are shaded, and just one paper [Shin 2020] cited BERT [Devlin 2018], the most widely used TLM, but only as a suggestion for future work. *No papers published or forthcoming in any of these journals at the time of this literature review had utilized CWRs, DTL or TLMs.*

For brevity, we do not discuss the details of the results of the IS literature review reported in Table 1, but we will reference different elements of this table in the remainder of the document. In some cases, we reference specific studies from this section to demonstrate how TLMs can be used to improve on this work. However, our review of IS literature is not limited to just these studies, and in the next section we also identify more existing research, from IS and other domains, that did not utilize text mining, but which still stands to benefit from TLMs.

The following section focuses on the most significant ways that TLMs and CWRs could impact IS research. While IS research examples are used in this, more detailed examples are described in Appendix A: Summaries & Recommendations. In this supplementary material, we include a summary of each paper from the IS literature review along with recommendations for how, or how not, the work might benefit from using TLMs or CWRs. We feel that this also can help readers to better understand how these new techniques are poised to impact IS research. This appendix was not included in the main text due to length, but we strongly feel that it is a major contribution of the paper; consequently, we recommend interested readers consider it to be primary content.

While Table 1 reports the results of the literature review, and supports the content presented in the remainder of the paper, it alone does not provide a valuable synthesis for IS researchers applying text mining and NLP in their research. A critical element of research involving text data is identifying appropriate techniques for different types of data. Table 2 shows the methods from our literature review that are most used for various types of text analytics employed in IS research. It can be seen that all methods are used for social media data, and that nearly all techniques are also used for reviews. However, for other documents (e.g., government documents, financial documents) and for text data collected from apps, only techniques like topic modeling and feature extraction (or word representation models) have been utilized. However, for more generic internet data that does not involve reviews or social elements, all techniques have been used, except for sentiment analysis.

In the following section we discuss how TLMs, CWRs and DTL can be used for each of the five classes of text mining techniques that we cover in Table 2. While we also cover more speculative applications for IS researchers, we feel that simply applying these new techniques in manners consistent with the most common applications from the existing literature offer the most promising opportunities for IS researchers.

4 IMPLICATIONS FOR RESEARCH AND PRACTICE

TLMs will enable researchers and practitioners to leverage the broad NLP power of transformers for task-specific applications through DTL via fine-tuning and through the development of custom

Table 1. Analysis of Text Mining Research in IS Journals

	Study	Sentiment Analysis	Topic Clustering	Feature Extraction	Entity Extraction	Text Classification	Emotion Detection	Other	Predictive Model	Reviews	Social Media	Novel Method	Text Mining Focus	Robustness Check	Foreign Language
Information Systems Research	Abbasi et al. 2019	✓					✓		✓		✓	✓			
	Adamopoulos et al. 2018		✓	✓			✓		✓		✓				
	Blohm et al. 2016							✓							
	Chen, K. et al. 2020	✓			✓			✓	✓			✓	✓		
	Chen, W. et al. 2019	✓								✓					
	Chung et al. 2020	✓									✓		✓		✓
	Heimback & Hinz 2018	✓		✓							✓				
	Hwang et al. 2019		✓								✓				
	Khern-am-naui et al. 2018	✓	✓	✓				✓		✓			✓		✓
	Lappas et al. 2016	✓		✓						✓		✓			
	Lee et al. 2020		✓	✓								✓	✓		
	Liu, X. et al. 2020		✓	✓		✓			✓		✓	✓	✓		
	Liu, Y. et al. 2020		✓	✓				✓	✓			✓	✓		
	Mejia et al. 2019					✓		✓	✓	✓		✓	✓		
	Mousavi & Gu 2019	✓							✓		✓				
	Mousavi et al. 2020	✓							✓		✓		✓		
	Pan et al. 2019							✓	✓			✓	✓		
	Song et al. 2019	✓		✓		✓		✓	✓		✓	✓	✓		✓
	Wang, Q. et al. 2018		✓						✓			✓			
	Yang, Mo. et al. 2018	✓				✓				✓					
Journal of Management Information Systems	Chen, L. et al. 2019	✓	✓	✓	✓	✓	✓	✓	✓		✓	✓	✓		
	Dong et al. 2018	✓	✓	✓		✓	✓	✓	✓		✓	✓	✓	✓	
	Ghiassi et al. 2016	✓									✓	✓			
	Huang, J. et al. 2017	✓							✓		✓		✓		
	Jha et al. 2019	✓						✓		✓				✓	
	Kumar et al. 2019									✓			✓		
	Li et al. 2016	✓	✓	✓					✓				✓		✓
	Mai et al. 2018	✓		✓					✓		✓				
	Saifee et al. 2019	✓								✓					
	Samtani et al. 2017		✓								✓	✓			
	Siering et al. 2016		✓			✓	✓		✓			✓	✓		
	Shi, D. et al. 2017					✓		✓	✓			✓	✓		
	Van Osch et al. 2018			✓		✓					✓				
	Velichety et al. 2019	✓		✓	✓				✓			✓			
	Wang, Z. et al. 2020		✓	✓					✓			✓	✓		
	Yoo et al. 2019							✓			✓				
	Zhang, D. et al. 2016			✓		✓	✓	✓	✓	✓					
	Zhou et al. 2018			✓			✓	✓		✓					
MIS Quarterly	Abbasi et al. 2018	✓		✓	✓	✓		✓	✓		✓	✓	✓		
	Arazy et al. 2020			✓				✓							
	Bapna et al. 2019		✓								✓				
	Benjamin et al. 2019				✓						✓	✓			✓
	Chau et al. 2020	✓		✓		✓			✓		✓	✓	✓		
	Deng et al. 2018	✓					✓				✓		✓		
	Gong et al. 2017		✓	✓									✓		
	Huang, K.Y. et al. 2019					✓					✓		✓		
	Huang, N. et al. 2016	✓		✓			✓	✓		✓					
	Li, J. et al. 2020		✓	✓	✓	✓		✓				✓	✓		
	Liu, X. et al. 2019				✓	✓					✓	✓		✓	
	Rhue & Sundararajan 2019	✓		✓						✓	✓				
	Shi, Z. et al. 2016		✓	✓				✓				✓			
	Shin, D. et al. 2019		✓	✓				✓	✓		✓	✓			
	Wu, J. et al. 2019	✓						✓	✓	✓		✓	✓		
	Yue et al. 2019		✓			✓		✓	✓		✓				✓
	Zhang, K. et al 2016	✓					✓		✓		✓				
	Zhang, W. & Ram 2020.			✓		✓	✓				✓				
Key	✓ = less than full utilization	✓ = full utilization							Using word representation						

Table 2. Text Mining Techniques and Their Applications in IS Research

Text Data Source	Methods				
	Sentiment Analysis	Emotion Detection	Text Classification	Topic Modeling	Word Vector Models
Reviews	✓	✓	✓		✓
Social Media	✓	✓	✓	✓	✓
Other Documents				✓	✓
Apps				✓	✓
Other Web Content		✓	✓	✓	✓
Key	✓ = used in IS literature		✓ = used feature extraction		

models with rich CWRs. In doing so, they offer opportunities to improve insights from existing research and they will open the door to powerful NLP-driven analyses across a wide variety of new domains and applications. Furthermore, they will enable fundamental changes in the nature of human-computer interaction by enabling the widespread use of LUIs. While it is not possible to anticipate the upper bound of the technical capabilities that TLMs will unlock, much less the ways in which TLMs will impact organizations and society, we still attempt to outline those ways that seem plausible based on the bodies of literature we have reviewed in this survey.

In the previous sections we have only considered the technical elements of TLMs and the existing body of IS literature that utilized text mining. In this section we further explore recent literature regarding TLMs, but with a focus on their applications in both research and practice. We do this in two phases: (1) by examining their potential for use in traditional text mining/NLP tasks and applications, and (2) by examining their potential for use on NLP tasks for those that have not previously been practically useful due to insufficient performance. Through each of these phases we consider tasks and applications through their standard typification in the text mining and NLP literature. Throughout this process we consider implications of our discussion on both research and practice, especially in areas where they can be used to enhance or broaden the body of existing IS research.¹⁶

4.1 Enhancing Existing Text Mining and NLP Applications

4.1.1 *Sentiment Analysis.* Sentiment analysis is the most widely used text analytics technique in recent IS literature (see Table 1) with a broad range of applications including ecommerce, market intelligence, social media analytics, government, politics, security and public safety [Chen 2012]. TLMs have already been used to improve upon the state-of-the-art (SOTA) performance on benchmarks for the major classes of sentiment analysis: aspect-based sentiment analysis, fine-grained sentiment analysis, targeted sentiment analysis and emotion detection [Phan 2020; Cheang 2020; Naseem 2020; Zhong 2019]. However, while we do feel strongly about the potential for TLMs, we also recognize sentiment analysis is a complex and well-established field and we do not intend to suggest that TLMs can nontrivially be used to improve on all existing applications. A comprehensive discussion of the implications of TLMs on sentiment analysis in IS research is beyond the scope of our study, but in this subsection we highlight the ways in which we feel that these techniques can be applied to benefit researchers and practitioners. Due to the prevalence of sentiment

¹⁶We note the caveat that, while many of the studies cited here produce state-of-the-art results, their results may not yet enable new forms of research as we suggest. However, we feel our suggestions are prescient and justified due to the rapid pace of progress, and due to the fact that most of the studies cited utilize BERT [Devlin 2019], which only represents the baseline in the SuperGLUE [Wang 2019a] benchmark.

analysis in IS research, we do not provide many explicit examples but rather focus on the potential of the latest developments.

For general applications, XLNet [Yang 2019] has demonstrated SOTA performance for the largest variety of benchmarks and is the best suited TLM for fine-tuning tasks involving task-specific sentiment analysis. More recently, the TLM SentiLARE has demonstrated strong all-around performance in sentiment analysis tasks by incorporating linguistic knowledge from SentiWordNet (Ke 2020). TLMs and CWRs have also been used to outperform LSTM and SVM-based methods on investor sentiment analysis [Li 2021] and to achieve SOTA results on targeted, domain-specific datasets such as airline industry Twitter data [Naseem 2020]. Implementations of TLMs such as these have strong implications for IS researchers utilizing sentiment analysis on domain-specific, targeted topics.

Some have gone as far as to suggest that BERT [Devlin 2019] be used as the standard baseline for comparing future progress [Li 2019]. We believe that this is reasonable, and suggest that widely used TLMs (e.g., BERT) should be used as baselines for comparing all relevant novel text mining or NLP methods moving forward. We further suggest that due to concerns about the impact of mismeasurement and misclassification error from extracted data mining features on the validity of IS research [Yang 2018] those who choose to use alternative sentiment analysis techniques as input features for statistical models offer better justifications for their selected methods, including explanations for why fine-tuned TLM models were not used or including comparisons to more advanced TLM-based models.¹⁷

Aspect-based sentiment analysis (ABSA) has great potential for business applications, such as for understanding online reviews in a finer-grained manner [Huang 2020], but our literature review indicates that it has not been applied in research published in premier IS journals yet. However, TLMs may make this easier in future work. DomBERT has been proposed to do just this (and more) by helping to train domain-specific TLMs with minimal resources, and it has shown promising results on ABSA tasks [Xu 2020]. Other recent work on improving analysis of online reviews has advanced the SOTA on widely used ABSA benchmarks of online reviews [Phan 2020], and we feel that using TLMs for ABSA is an excellent avenue for future IS research.

Due to the value of DTL, TLMs make targeted sentiment analysis a particularly easy area for improving upon the analyses of existing IS research when large amounts of data are available. Even in cases where labeled training datasets are not available, crowd sourcing options such as the Amazon Mechanical Turk make the labeling of a modest number of documents¹⁸ reasonable. This is particularly useful for cases involving text data, like Tweets, that do not conform to standard syntactic rules, or cases when specific topics are of interest. For an IS example, Ghiassi et al. [2017] developed a custom sentiment model using feature engineering, vectorization and a trained classifier. However, TLMs remove the text mining knowledge necessary for developing an advanced model like this and make it more straightforward to achieve optimal performance on business related datasets, such as that used by Ghiassi et al., with only data collection, cleansing and sentiment scoring. For the same reasons, TLMs offer benefits to practitioners, where marginal improvements in the quality of input features and the quality of results can have a more tangible and valuable effect than in research by increasing revenue, sales or profits.

4.1.2 Emotion Detection. Emotion detection is a form of sentiment analysis that we feel is worth mentioning separately due to its relevance to IS research. From text alone, it can be particularly

¹⁷Such as fine-tuned XLNet [Yang 2019] or SentiLARE [Ke 2020]. See Phan and Ogunbona [2020] for an example.

¹⁸A number of documents on the order of magnitude from 1,000 to 10,000 is often sufficient for fine-tuning pretrained TLMs.

challenging due to the absence of knowledge about the target's gestures or facial expressions [Chatterjee 2019a]. Progress on this task has proved to be more challenging than some of the other tasks discussed in this section where BERT-based models have easily improved upon SOTA results. However, progress has still been made, for example, by fine-tuning on evaluation datasets using TLMs that account for commonsense reasoning by incorporating a commonsense knowledge base and an emoticon lexicon during pretraining¹⁹ [Zhong 2019]. Commonsense knowledge has also been employed more recently by using pretrained CWRs to incorporate different commonsense elements such as mental states and causal relations to learn interactions between interlocutors in dialogue, achieving SOTA performance on four conversational emotion benchmarks [Ghosal 2020]. Emotion detection is more challenging than other tasks that we have focused on, but these early results offer strong evidence for the ability of combining CWRs and TLMs with other techniques (e.g., knowledge graphs) to outperform existing models on such challenging tasks.

While more complex, these solutions make progress on a topic that is particularly valuable in IS research and business more broadly. A number of studies published in elite IS journals involve emotion, however, they often involve designed experiments [Liang 2019] or qualitative and mixed-methods approaches [Salo 2020]. Consequently, we feel that methodological research involving TLMs for emotion detection is a topic that is well-suited for IS researchers and should be prioritized due to its potential impact in IS research and beyond. We further expect that emotion detection could have an impact on other areas of business research such as marketing²⁰ and finance.²¹

In our IS literature review, Chau et al. [2020] demonstrated a novel model utilizing a text mining driven classifier in tandem with a rule-based classifier to identify at-risk individuals exhibiting emotional distress. This is a novel and important application of text analytics in IS research, yet, in light of the literature reviewed here there is much room for improvement²², and it would be interesting to see the methods discussed in this section used for future work along these lines. The data used by Chau et al. was in Chinese, so some of the techniques suggested above may not have been viable alternatives, but multilingual TLMs, discussed later in this subsection, offer new solutions for this as well.

4.1.3 Text Classification. Text classification is an essential technique of text mining that has numerous applications in organizations. While it was not as widely used in our survey as sentiment analysis or feature extraction, it was commonly used in combination with other text mining techniques and is one of the techniques which stands to improve most dramatically from fine-tuning TLMs. This is underscored by the fact that, in their seminal paper on transfer learning for language models, Howard and Ruder [2018] focused on six text classification tasks for

¹⁹This example differs from the majority in that it involved the specialized pretraining of a TLM as well as task-specific fine-tuning (as opposed to using an out-of-the-box TLM pretrained on a large, generic corpus).

²⁰As an example, from marketing, Rocklage and Fazio [2020] examine the effects of emotion in online reviews using a lexicon-based emotion analysis technique. We feel that this could likely benefit from more fine-grained analysis using either ABSA-based methods or some of the more complex commonsense-based emotion detection techniques discussed above.

²¹In finance, studies focus on sentiment and emotion, but they do not use text mining techniques [Jiang 2019; Cortes 2016]. We feel that these new methods may be strong enough to lead to valuable insights which can aid in the development of new theoretical contributions, and we suggest that researchers and practitioners in these disciplines consider applying the methods discussed in this subsection for datasets available in their domain.

²²For one, it uses a lexicon-based method for feature extraction, which is not relevant enough for comparison on emotion recognition benchmarks in our literature review or in the foremost computing psychology journal [Chatterjee 2019b]. We feel that it would have been useful for a study so recent to have included a comparison to the current methods discussed here.

demonstrating the value of DTL in NLP. BERT [Devlin 2019], fine-tuned on domain-specific datasets, was quickly demonstrated to achieve SOTA performance for a variety of text and document classification tasks [Yao 2019; Sun 2019b]. One example that could be particularly useful for IS research is BERTweet, which is a BERT-based model pretrained on Twitter data that achieves SOTA performance on Twitter text classification as well as part-of-speech-tagging and named-entity recognition [Nguyen 2020]. Models like this, pretrained on domain-specific data, are quite common: SciBERT [Beltagy 2019] and COVID-Twitter-BERT [Müller 2020]. Such models can then be fine-tuned on task-specific data for further performance gains. Due to the improved performance they bring, it is likely that similar models could be very useful for numerous applications in IS research, other business domains and the social sciences more broadly. For an IS example, one could extend the work of Mejia et al. [2019] on classifying restaurant hygiene by unsupervised pretraining of a BERT-based model on bulk restaurant reviews, then fine-tuning for classification of “instances of hygiene violations.”

TLMs enable the creation of text classification models which previously required complex methods to be created with significantly less effort and expertise. Huang et al.’s [2020] study of support and companionship in virtual healthcare communities offers an excellent opportunity to use fine-tuning to improve model performance. BERT [Devlin 2019] could be fine-tuned via a Google Colab notebook²³ (and a powerful coprocessor²⁴) for free²⁵, as could smaller T5 models [Raffel 2020]. However, Colab offers a good opportunity to debug T5, and Huang et al.’s study offers a good opportunity to use T5’s multitask capability. As another example, Kraus and Feuerriegel [2017] developed a Bi-LSTM model for predicting a firm’s market performance based on financial disclosures, but BERT could be fine-tuned on the same data using Colab and a few dozen lines of code to improve performance (see Wolf 2019).

Training a TLM to classify the data from Kraus and Feuerriegel [2017] would work by simply inputting entire documents because the model would simply output a class, but not all documents are short enough to fit in the context window of TLMs.²⁶ Innovative models such as DocBERT [Adhikari 2019] and the Longformer [Beltagy 2020] have achieved SOTA results on various document classification tasks, as well as other document related tasks, and could be useful for longer documents like internal reports, legal documents, newspaper and magazine articles or longer Wikipedia articles. Moreover, recent modifications to the original transformer architecture such as the reformer [Kitaev 2020] suggest that larger context windows will be a feature of TLMs in the near future.

4.1.4 Topic Modeling. Topic modelling is widely used in IS research, as indicated by our survey, and the most widely used technique is **latent Dirichlet allocation (LDA)** [Blei 2003]. However, TLMs have also performed well in these areas and BERT [Devlin 2019] has been shown to improve upon the SOTA when applied to specific use cases such as argument [Reimers 2019] and document clustering [Park 2019]. Moreover, contextual document embeddings from TLMs have been shown to improve topic coherence [Bianchi 2020]. However, overall it is unclear whether BERT-based CWR clustering improves on LDA enough to make a difference [Sia 2020], but the results from Sia et al. suggest that larger TLM CWRs such as those from RoBERTa [Liu 2019c], XLNet [Yang 2019] or T5 [Raffel 2020] could be expected to outperform LDA. While there may be some uncertainty about

²³Google Colab notebooks (virtual Jupyter notebooks) can be used to run simple deep learning models directly through the browser, and we feel that these are well suited for most applications of TLMs and CWRs suggested in this study. Official instructional notebooks exist for all of the most widely used models, and 3rd party notebooks exist for many other models.

²⁴This will either be a SOTA graphics processing unit (GPU) or one of Google’s proprietary tensor processing units (TPUs).

²⁵Colab is free but has usage limits. However, Colab Pro, for \$10 per month has no limits and more generous TPU allocation.

²⁶There is a limit to the number of tokens that can be input (e.g., 512 for BERT) [Devlin 2019], though this context window is larger for larger models such as T5 (e.g., up to 2,048) [Raffel 2020].

using CWRs for clustering, Hoyle et al. [2020] have demonstrated that TLM-based techniques can be used to obtain SOTA topic coherence. They do this not by using CWRs or TLMs directly for topic modeling, but by using their **BERT-based Autoencoder Teacher (BAT)** approach in tandem with SOTA topic modeling methods. Thus, this is another case in which TLM-based methods should begin to be used as default methods. This can have important implications for IS research because improved input features can significantly impact the statistical validity of IS research results [Yang 2018].

4.1.5 Word Representation Models. Word representation models are commonly used in IS research, particularly when feature extraction is necessary. While such techniques do not outperform CWRs, they are still able to perform relatively well on tasks with plentiful data and simple language [Arora 2020], but our review has indicated that this is not always the case for IS. Thus, we see numerous studies as being able to benefit from the improvements offered by CWRs. For example, Arazy et al. [2020] focus on the evolution of digital artifacts (i.e., wiki articles) over time by tracking trajectories in a feature space, and, because the authors do not use text mining, they explicitly suggest the use of word representations would benefit future work.²⁷ As another example we consider Wang et al. [2020] who extract soft semantic factor characteristics from descriptive loan texts, but the semantic similarities between words and loan texts could be more easily and effectively captured in a latent feature space using CWRs. Numerous other studies utilize feature extraction, some even using neural word representations, and many stand to gain from using more advanced CWRs (see Appendix A for concrete suggestions on the papers from the IS literature review).

Other models in the IS literature have used alternative techniques for feature extraction to develop novel distributional representations of text [Shi 2016], and we feel that some of these models offer good opportunities for using CWRs. For example, Shi et al. used LDA [Blei 2003] for feature extraction to represent aspects of firms' business to evaluate firms' relative "business proximity." Lee et al. [2020] also use LDA to create a novel "app similarity measure." Work such as this is well suited for CWRs which can be crafted in a custom fashion to create novel measures of documents' semantic similarity [Gyawali 2020]. In general, LDA is widely used in the IS research literature for extracting features [Gong 2018; Shin 2020; Liu 2020b], but even if dimensionality reduction is necessary for using the features in statistical models, we agree with Shin et al. that CWRs can provide richer representations.

4.2 Beyond Existing Applications

While CWRs and TLMs have significant implications for improving and furthering existing IS research, we believe that their most interesting applications for IS research are in their advanced and novel applications. In this subsection we discuss these emerging topics.

4.2.1 Regression. Regression is one emerging application for which little previous work has been conducted in NLP. One good example of using NLP for regression was by Kraus and Feuerriegel [2017] who used financial disclosures to predict firms' subsequent performance in financial markets, but their model required a very specialized LSTM model. However, it is possible to simply fine-tune language models for regression by posing regression problems as text-to-text tasks [Raffel 2020]. While, this is still an emerging research area, it has been demonstrated for applications such as table retrieval [Chen 2021] and to predict brain activity as measured by fMRI based on the text being read [Schwartz 2019]. One practical example of regression on text data is that of

²⁷Specifically, they suggest that "the feature space could be represented through more sophisticated text processing methods and more advanced knowledge representations ... the patterns observed here serve as a lower bound" [Arazy 2020].

automated essay scoring such as for standardized tests. Yang et al. [2020] find that simply fine-tuning on BERT [Devlin 2019] is not enough, but that extracting CWRs from BERT and training a fully-connected neural network on multiple losses improves on the SOTA performance by almost 3%. We feel that this example offers promise for many practical business tasks, as well as for numerous uses in IS research.

4.2.2 Multilingual Analytics. Machine translation can be useful in business intelligence and business analytics applications when organizations need to analyze or monitor either static or streaming text data in multiple languages [Moreno 2016]. While it is more commonly thought to be an independent research area within NLP, like speech recognition, progress on the related topic of multilingual language models does have significant implications for IS research. Machine translation has more applications in practice than research, and interested readers are encouraged to review recent high-level overviews [Hao 2019]. Here, our discussion focuses broadly on multilingual capabilities of TLMs and their applications in both research and practice.

Multilingual TLMs were introduced in the previous section and models like XLM-R [Conneau 2020] result in significant improvements for a wide variety of cross-lingual transfer tasks. What is most interesting about the results from Conneau et al. is that they suggest these gains may be possible without sacrificing monolingual performance. It may not be obvious how this will impact IS research, but there is a digital language divide between dominant languages [Young 2015], and as information technology has proliferated over time this divide has had a significant impact on their adoption and applications across cultures. Consequently, multilingual TLMs enable a powerful tool to examine this using deep learning analytics. For an IS example, George et al. [2018] evaluated the effect of communication media and culture on deception detection by conducting an experiment which showed that different combinations of media and cultural effects affected deception detection accuracy. Multilingual language models offer the ability to conduct research in this vein without the burden of conducting an experiment with groups across three different languages, a burden that is likely prohibitive to most IS researchers. Our survey of IS literature found social media and online reviews to be the primary applications of text mining in IS research. Simply considering social media, and the ability to apply TLMs for analysis of behavior across cultures, one can look at recent work in leading human-computer interaction journals [Wang 2019b; Cho 2018] and foresee the strong research potential here. Thus, we anticipate that multilingual TLMs will open doors to numerous new research directions for IS researchers (e.g., Ebrahimi et al. [2021]).

Yet, these models' value is not limited strictly to cultural comparisons and can be applied directly to improve insights from existing IS research. The work of Chau et al. [2020] mentioned earlier could benefit from using multilingual representations to replace older lexicon-based methods of feature extraction. This hints at the possibility of being able to conduct IS research on non-English datasets without the need for fluency in the language of focus. If possible, this would open up a wide variety of foreign language datasets to IS researchers.

4.2.3 Language Generation. Language generation has been a topic of interest in the NLP community for over a decade, and it is such a significant topic with respect to TLMs that we make a distinction between standard TLMs and generative language models. GPT-2 [Radford 2019] was the first generative language model to really demonstrate shockingly impressive language generation results. It was followed by T5 [Raffel 2020] and most recently by GPT-3 [Brown 2020], which each demonstrated shocking gains.

Significant effort is going into improving reliability and ease of generating samples that are more human-like [Keskar 2019] or less biased [Huang 2020; Ma 2020] while others are focusing on applying TLMs to more immediately practical applications, such as chatbots [Roller 2020b]. We previously mentioned chatbots that were closing in on human-level performance for open domain

conversation [Adiwardana 2020; Roller 2020a]. We expect language generation to be inextricably related to the future of IS research in a very significant way given its potential to fundamentally change human-computer interaction. The remainder of this section focuses on different applications of language generating systems with implications for future IS research such as for document summarization, question and answering, automated report generation and language user interfaces.

4.2.4 Document Summarization. Document summarization is a task that has the potential to be very valuable for business intelligence and business analytics applications. While document summarization is still a very challenging task [Kryściński 2019], TLMs are showing promise in this area, and have even successfully been able to use recursive summarization schemes for summarizing entire novels [Wu 2021]. Generally, document summarization is classified as being one of two types: extractive or abstractive. Extractive summarization involves identifying and concatenating extracts from the document into a summary. Improvements for this using TLMs are straightforward for specialized applications because existing pre-trained models can simply be fine-tuned on domain-specific datasets [Gu 2019]. However, abstractive summarization is more challenging, yet, despite this, more complex TLMs have been able to achieve SOTA performance when trained directly on task-specific datasets [Duan 2019]. Researchers have begun using unified frameworks for multitask models capable of both abstractive and extractive summarization [Chen 2019], leading to SOTA on benchmarks for both extractive and abstractive summarizations [Liu 2019b] and multi-document summarization [Jin 2020]. Abstractive summarization is more valuable in the long run, and recent work on this task has concluded that TLMs and generative language models are able to generate more informative, coherent, faithful and factual summaries [Maynez 2020].

The potential applications of summarization for business intelligence systems are wide-ranging. For one, if we extend summarization to full report generating systems, we can envision how such systems could leverage industry reports, news articles and social media to power business intelligence systems with real-time understanding of complex market behavior in the form of an intelligent dashboard. Summarizations could also be used for reducing reading time on emails or other long documents or reports produced by employees at all levels of the organization. The ability to highlight the key points in a document may even be more beneficial in this aspect. Exciting new work from OpenAI has shown significant improvements in summary quality by using human feedback to train summarization models [Stiennon 2020] and these results suggest that practical use of summarization systems may not be far away.

Multi-document summarization [Lu 2020] and extreme summarization [Narayan 2020] have become popular topics as well, and ones with significant implications for practical applications in highly specialized domains (e.g., science, finance, etc.). Extreme summarization refers to summarizing highly technical documents, such as scientific papers, with a single sentence. This could also be very useful for summarizing financial statements or legal documents. Documents of this sort are often large in number, and query focused multi-document summarization that is effective for a range from coarse-to-fine estimation [Xu 2020] could be extremely useful in future business intelligence systems for these domains.

Another potentially very useful application of summarization would be practical cross-lingual summarization, which would use a multilingual TLM to generate a summary in one language from a text written in another language. TLMs have been used for this, but the relative performance was not possible to determine [Zhu 2019]. Exciting new work shows continued progress on this task [Cao 2020], and a new multilingual summary dataset [Scialom 2020] and benchmark [Ladhak 2020] suggest that we can expect more work on this topic in the future. Similar to our discussion of multilingual TLMs earlier, the applications discussed in this subsection offer numerous opportunities for IS researchers and open the door to novel research questions.

4.2.5 Question Answering. **Question answering (QA)** systems that are effective for domain-specific applications have tremendous potential for business intelligence. Such systems could fundamentally change the nature of decision support for any application with enough data for fine-tuning. Impressively, systems have been able to score an A on a standardized New York 8th grade science exam and a B – an 83 – on the same 12th grade exam [Clark 2019]. While this sort of generality is not necessary for practical applications, it effectively demonstrates how powerful QA systems from TLMs can be. For many practical business applications, a high school graduate that can make an 83 on the most difficult standardized high school level science exam can likely read documents and be able to generate answers that suffice for a wide range of data and applications relevant to organizations. Many tasks that standard college graduates do in white collar jobs do not require the full use of their faculties and education.

Such powerful QA systems, particularly when it is possible to fine-tune them for customized experiments, have the potential for valuable new directions in IS research. For example, we can consider QA systems that are easily fine-tuned on domain-specific datasets. This has been a desirable goal for many years, especially since IBM's Watson, but it has not materialized as many had originally anticipated. Yet, given the rapid progress of TLMs, we can expect such systems to become practical in the near future. Xu and Lapata [2020] discuss adapting recent QA methods to improve query focused multi-document summarization, and such systems have the potential to transform strategic decision making in organizations and dramatically impact the nature of white-collar labor. If we consider not just multi-document QA, or data warehouse QA, but QA based on an entire organization's archived text data, we can begin to understand this potential. Yet, however transformative these technologies may be, it is likely that these advances will first lead to the augmentation of human jobs rather than the replacement of them [Morgan 2019], and it falls on IS researchers to develop an understanding of how this augmentation of occupations will impact organizations and the future of white-collar work.

Recent work on a knowledge-intensive generative language model from Facebook – **retrieval-augmented generation (RAG)** – demonstrated SOTA performance for three widely used, general QA tasks [Lewis 2020c]. In the same month, another QA oriented generative language model from the Allen Institute, called UnifiedQA, demonstrated strong performance without fine-tuning, and was able to achieve SOTA performance on 10 factoid and commonsense QA benchmarks [Khashabi 2020]. However, while all of this may seem impractical due to the lack of labeled datasets, there are strong, user-friendly extractive QA systems that can be fine-tuned on large, unlabeled domain-specific datasets [Dibia 2020] which could be used for IS research now and which could offer guidance for future research as business-related question answer datasets are created and as systems grow more capable.²⁸

Generally, work on reading comprehension is closely related to QA systems. Thus, it should come as no surprise that CWRs were already rivaling SOTA performance in related tasks in 2018 [Salant 2018]. After its release, BERT [Devlin 2019] soon achieved new SOTA performance on multiple benchmarks in multiple choice reading comprehension tasks [Zhang 2019]. Based on the prevalence of online reviews that our survey illuminated in existing IS research, the new idea of *review reading comprehension* proposed by Xu et al. [2019], for including a QA system on top of a large repository of ecommerce reviews, may offer some further insight into the potential for fine-tuning multitask models on domain-specific datasets. Their system targeted customers, but similar systems could be developed for other applications in organizations such as for analysts working to increase revenue or to improve customer satisfaction.

²⁸All of these models are available open-source: RAG at https://huggingface.co/transformers/model_doc/rag.html; UnifiedQA at <https://huggingface.co/allenai/unifiedqa-t5-large>; NeuralQA at <https://github.com/victordibia/neuralqa>.

4.2.6 Language User Interfaces. **Language user interfaces (LUIs)** have long been anticipated to become a widely used modality of human-computer interaction [Brennan 1991]. While LUIs still play only a limited role in our daily interactions with computers, recent progress in TLMs raises the possibility that LUIs will become practical and widespread in the near future. We envision practical LUIs to be powerful systems that are used to enhance human capabilities through human-computer interaction [de Vries 2020]. In the following paragraph we will briefly discuss some possibilities for practical applications of LUIs.²⁹

We define an LUI as an intelligent system that is goal oriented to substantially enhance economically valued human capabilities through an interface that is optimally controlled with natural language. We are particularly interested in LUIs that are practical in the sense that they can assist humans in tasks of nontrivial economic utility. Some obvious examples of LUIs are personal assistants, assistants for the impaired or customer support assistants. While many call centers already use automation and there are widely used personal assistants like Google Assistant and Siri, their economic utility is relatively limited. Perhaps this is truer for the personal assistants than for the call centers but call center automation has been gradually increasing for decades. Many tasks are repetitive, like navigating information systems, and they do not require strong language understanding or interaction. Thus, such systems do not meet the criteria of being optimally controlled through natural language. Furthermore, while we do not feel any of these existing example systems meet the criterion of enhancing economically valued human capabilities, we feel that TLMs are currently poised to usher in dramatic progress on it.

More powerful LUIs that we foresee include navigation agents for automobiles or flying vehicles, interactive domestic appliances and domestic robots. Furthermore, directly related to organizations' productivity, we envision agentive business intelligence systems that are able to offer powerful capabilities such as those mentioned earlier in this subsection like summarization or QA capabilities, but which also leverage reinforcement learning to tailor their functionality to a specific user. Such systems would truly transform the nature of business intelligence and decision support systems, and it is critical for IS researchers to begin understanding how these systems will change organizations and society in the years to come because their rise may come quickly [Gruetzmacher 2020].

While this is primarily a topic for future research, it is possible for eager IS researchers to begin work on these problems at present. We have included links (in footnotes) in this subsection to open-source code that could be used to these ends. Further, Roller et al. [2020a] demonstrated the best performing chatbot³⁰ to date in their Blender Bot while also releasing the code open-source as well as the 9.4 billion parameter pretrained model.³¹ We believe that this alone unlocks a wide range of novel IS research directions, particularly if the model is fine-tuned for specific tasks and evaluated empirically. Other recent work on ToD-BERT [Wu 2020] for task-oriented dialogue offers another tool³² for conducting preparadigmatic research in this new domain. We feel that such research is important because children are already accustomed to LUIs like Alexa and Siri in their homes and phones and are beginning to expect devices to respond to verbal commands; we anticipate that in the coming decades, when entering the workforce they will expect language-enabled support in the workplace.

4.2.7 Few-Shot Learning. Few-shot learning is something that we feel will be closely related to LUIs, but its potential impact has strong enough potential to garner a brief but independent

²⁹ A full discussion of LUIs is beyond the scope of this paper, but interested readers are referred to (authors' working paper).

³⁰ While there has been dramatic and practical progress on open domain chatbots recently, a full discussion of chatbots is beyond the scope of this study. (This topic will be covered in greater detail in forthcoming work from the authors on LUIs.)

³¹ This can be found at: <https://parl.ai/projects/recipes/>.

³² The code can be found at: <https://github.com/jasonwu0731/ToD-BERT>.

discussion. The strong performance of GPT-3 on certain tasks such as QA via zero- one- and few-shot learning suggests the possibility of novel LUIs, and we feel that this is a topic that also falls to IS researchers to explore. It is beyond the scope of this study to illuminate in detail the potential for few-shot learning in IS research, but we suggest considering the following. Generative language models like GPT-3 take a text prompt at the time of inference, and, in the case of GPT-3, this prompt can be long and involve sequential tasks such as questions followed by answers. Performance from GPT-3 on tasks demonstrated in this manner is particularly strong, as we discussed in an earlier section.

GPT-3 also performs well on prompts that give a context and ask for the model to fill in the blank, to complete the sentence or even to generate an essay based on the prompt. Thus, it is easy to see how continued research and increasing the scale of powerful generative language models like GPT-3 can lead to very useful systems capable of report generation, summarization and other valuable tasks if trained with a larger context window (and at a higher computational cost). However, what is not as obvious is the value of direct user interface with a system capable of learning complex contexts such as QA or mathematical operations. It is likely that there are novel ways of interacting with such interfaces that can create value in ways that are difficult to imagine *a priori*. For example, one startup is using GPT-3 exclusively to improve inbox productivity by generating detailed emails from short prompts.³³ They do this through a novel notion of LUI wherein the user does not have to reply in complete sentences, they only provide the information necessary for the response and the model generates an email response in context³⁴ with the correct information. We feel this is an appropriate and urgent topic for IS research.³⁵

5 SUMMARY OF IMPLICATIONS FOR IS RESEARCH

We first surveyed the recent progress in NLP that has led to SOTA performance on a wide range of tasks using TLMs. Next, we discussed and reviewed IS research that has used existing text mining and NLP techniques, a substantial portion of which could be improved by using CWRs or TLMs. We then discussed some of these possible improvements in the next section³⁶ as well as a number of possible avenues for new and novel IS research stemming from TLMs. In this section we summarize our findings and their implications for IS research.

TLMs have a handful of distinct and noteworthy advantages over standard text mining techniques. Foremost, they are able to achieve SOTA results on a wide variety of text mining and NLP tasks as long as a modestly sized dataset is available for training. However, their value is not limited to labeled datasets and fine-tuning; they can be used to generate rich CWRs which can be used to extract features for building custom models in combination with a variety of machine learning or statistical methods. Based on how often feature extraction has been used for text analytics in the recent IS literature, we feel that this alone can have a significant impact on future IS research (e.g., Samtani et al. [2021]).

The remainder of this section focuses on four distinct topics. First, we review the implications of TLMs which is followed by a discussion of new opportunities and LUIs. We next offer some brief suggestions for reviewers and editors when considering submissions involving novel methodological contributions for text analytics and finally, we discuss some implications of further TLM scaling and continued progress in grounded semantics.

³³Demonstrations can be seen at <https://www.OthersideAI.com>.

³⁴We imagine this fails on emails over 2,000 words because the entire email must be given to GPT-3 as context. It is also likely that the firm has fine-tuned GPT-3 for this task, which is very costly but would quickly bring a novel product to market.

³⁵OpenAI has recently begun licensing API access to GPT-3 through Microsoft.

³⁶A discussion for each study in the IS portion of the survey is included in tables A1-A3 in Appendix A.

5.1 Transformer Language Models (TLMs)

From sentiment analysis to emotion detection to text classification to regression to cross-lingual analysis, TLMs promise to have a significant positive impact on future IS research, notwithstanding the more advanced novel applications, and they can do so in a number of different ways. For improving existing work, they can either (1) be used to generate rich CWRs or (2) be used directly, through pretraining, through fine-tuning and DTL, or both. More exciting, they can (3) be used to extend existing IS research by enabling easier cross-cultural analyses. We describe these themes in this subsection and discuss other issues that may impact the future use of TLMs.

(1) When a modestly sized dataset is available, simple DTL and fine-tuning will often outperform all methods other than specialized TLMs or advanced models using CWRs and possibly LSTMs. This is important because using DTL to obtain such strong performance is significantly easier than the development of a custom LSTM model or the development of a custom TLM model or one that has not been pretrained. This should enable the wider use of fine-tuned TLM models for tasks such as sentiment analysis or text classification, thereby improving performance, even if only as a component of a more complex analysis. Because Google offers free, powerful Colab notebooks that include tutorials for fine-tuning a standard BERT model,³⁷ we feel it is reasonable to expect IS researchers to be able to do this with minimal machine learning expertise.³⁸

(2) CWRs generated from TLMs are superior for feature extraction from large datasets which are able to support high-dimensionality machine learning models. When this is not possible, CWRs can still be very valuable for feature extraction when coupled with dimensionality reduction and feature selection techniques. Due to the prevalence of feature extraction in our survey of IS literature, we feel that CWRs should be more widely used as it stands only to benefit IS research by reducing bias from mismeasurement [Yang 2018].

(3) Multilingual TLMs enable DTL to leverage multilingual representations for cross-lingual analytics. While this is still an emerging topic in NLP research, models such as XLM-R [Conneau 2020] are able to maintain monolingual performance while also being able to generate valuable representations for other languages. This can be very useful for IS research because it enables analysis of the effects of culture on social media behavior, technology acceptance, technology use, etc., all without having to design an experiment involving multiple languages. Moreover, it enables the use of SOTA methods when working on datasets involving foreign language data so that research in elite IS journals does not have to rely on older, more rudimentary methods [Chau 2018].

5.2 Novel Applications & LUIs

LUIs have been introduced as an incredibly promising area for IS research due to the recent progress in TLMs. While a full discussion of TLMs is beyond the scope of this paper, it is easy to see a path toward LUI research already emerging in the form of strong pretrained chatbots such as BlenderBot [Roller 2020a]. This chatbot is available open source³⁹, including the pretrained 9.4 billion parameter model, which enables IS researchers to begin working directly on LUIs.

Our literature review revealed few text analytics systems using design science, but LUIs will offer novel opportunities for using design science for artifact development and theorizing. Samtani et al.

³⁷See: https://colab.research.google.com/github/tensorflow/tpu/blob/master/tools/colab/bert_finetuning_with_cloud_tpus.ipynb.

³⁸Recall that pretraining schemes have a significant impact on downstream tasks: autoencoding models excel at discriminative tasks, autoregressive models excel at generative tasks, and sequence-to-sequence models attempt to balance performance between generative and discriminative tasks. Also recall that further fine-tuning and pre-finetuning can be utilized to enhance performance on many domain-specific tasks.

³⁹See footnote 24.

[2020] offer a strong template for such work, and we suggest that interested parties refer to it. We feel that this is a very strong potential area for research, and, because it is possible given existing technology, we suggest that interested IS researchers act fast to establish first mover advantage. We are eager and excited to see where research in this direction takes us.

5.3 Guidelines for Methodological Novelty Using TLMs and CWRs

TLMs offer a huge opportunity for researchers to improve upon previous SOTA results and to apply powerful NLP models to a wide variety of new applications which were previously not possible. With respect to improving upon SOTA, the ability of TLMs to do this is often related to the novelty and size of a training dataset rather than the novelty and methodological contributions of a technique. Thus, we encourage reviewers to be weary of papers employing TLMs which claim to make methodological contributions and to continue to seek theoretical contributions from novel studies using TLMs. However, this is not to say that methodological contributions cannot be made involving TLMs, but we suggest that it is necessary to compare results from proposed novel methods, like that of Chau et al. [2020], to static word representations as well as widely used CWRs. We further suggest that if TLMs or CWRs are used as a component of a proposed methodological contribution, the methodological contribution be made clear and robustly justified (e.g., specialized pretraining such as for Zhong et al. [2019]). The deep learning IS research template of Samtani et al. [2020] is also useful for this.

5.4 TLM Scaling and Grounded Semantics

AI practitioners anticipate a continued trend in the scaling of computational resources to continue to drive progress in AI research for the next decade [Gruetzmacher 2020]. Taken with recent research from OpenAI [Kaplan 2020; Brown 2020] this suggests that TLM progress will continue to improve dramatically, but that the costs of this increased performance will be non-trivial and may make research and operationalization of the most powerful TLMs quite costly, possibly even cost prohibitive. As noted, OpenAI has already begun licensing an API for the largest GPT-3 model through Microsoft; the prices are anticipated to be extreme for fine-tuning but more reasonable for few-shot learning. However, it is also difficult to anticipate how quickly progress to increase language model efficiency, such as adapters [Houlsby 2019], might progress and interact with the AI practitioner forecasts for scaling.

If others follow the API licensing model, it has the potential to dramatically impact the future use of TLMs in both positive and negative ways. Most obviously, it could put TLM research and operationalization out of reach for many academics and firms, at least for lower priority projects. Alternately, the high cost of the service demands a user interface that ensures users will not waste their time with the API and incentivizes the provider to make the product easy to use and maximally effective. This could significantly impact firms' adoption of TLMs as well as their use in research and is an interesting research question for future work.

Continued progress in grounded semantics could also have a dramatic impact on the performance and practicality of TLMs. We feel strongly that higher levels of grounding, such as embodiment and social [Bisk 2020], are certainly not necessary for language grounding to begin to start seeing practical applications. Again, it is difficult to anticipate how quickly progress may come, but it is likely that existing work, once refined, can have an impact on the use of TLMs in text analytics and for applications such as question answering or summarization.

6 CONCLUSIONS

In this work we reviewed two bodies of literature: (1) literature related to recent progress in NLP and (2) recent literature involving the application of text mining and NLP published in the top

IS journals. While some of the technologies we have discussed may mature over an extended period of time, it is important for IS researchers to keep up with the SOTA and to incorporate it into research without haste. This is true for all methodological progress, but it is particularly important for TLMs, CWRs, multilingual TLMs and LUIs as they have the potential to drive novel forms of IS research and substantially alter human labor and organizational processes for which text data is significant component. Even if such technologies are not mature, it is important for IS researchers to preemptively develop theory and methods for researchers and practitioners to use when the technologies do mature. We feel strongly that the IS research community, by more closely following progress in the NLP domain, can enhance the quality and value of their research contributions substantially. To these ends, we suggest the IS research community begin sponsoring workshops at the premier conferences (e.g., NeurIPS, ICLR, ICML, ACL and EMNLP⁴⁰) for business applications of these technologies.⁴¹

Overall, the literature and the ensuing discussion led us to conclude that transformer language models are poised to dramatically reshape the use of text analytics and NLP in IS research and practice. Moreover, by enabling technologies such as language user interfaces, they are likely to precipitate transformative change in organizations and in society. Taken together, these topics offer significant opportunities for future work in IS research and we look forward to seeing what the future holds.

ACKNOWLEDGMENTS

We thank Miles Brundage for comments on an earlier version of this manuscript. We also thank anonymous reviewers from the 2020 Winter Conference on Business Analytics for pointing out the need for an independent survey paper on this topic.

REFERENCES

- Ahmed Abbas, Yilu Zhou, Shasha Deng, and Pengzhu Zhang. 2018. Text analytics to support sense-making in social media: A language-action perspective. *MIS Quarterly* 42, 2 (2018), 427–464.
- Ahmed Abbasi, Jingjing Li, Donald Adjeroh, Marie Abate, and Wanhong Zheng. 2019. Don't mention it? Analyzing user-generated content signals for early adverse event warnings. *Information Systems Research* 30, 3 (2019), 1007–1028.
- Panagiotis Adamopoulos, Anindya Ghose, and Vilma Todri. 2018. The impact of user personality traits on word of mouth: Text-mining social media platforms. *Information Systems Research* 29, 3 (2018), 612–640.
- Ashutosh Adhikari, Ashyudh Ram, Raphael Tang, and Jimmy Lin. 2019. DocBERT: BERT for document classification. arXiv:1904.08398. Retrieved from <https://arxiv.org/abs/1904.08398>.
- Daniel Adiwardana, Minh-Thang Luong, David R. So, Jamie Hall, Noah Fiedel, Romal Thoppilan, Zi Yang, Apoorv Kulshreshtha, Gaurav Nemade, Yifeng Lu, and Quoc V. Le. 2020. Towards a human-like open-domain chatbot. arXiv:2001.09977. Retrieved from <https://arxiv.org/abs/2001.09977>.
- Armen Aghajanyan, Anchit Gupta, Akshat Shrivastava, Xilun Chen, Luke Zettlemoyer, and Sonal Gupta. 2021. Muppet: Massive multi-task representations with pre-finetuning. In *Proceedings of Empirical Methods in Natural Language Processing 2021. ACL*, 5799–5811.
- Ofer Arazy, Aron Lindberg, Mostafa Rezaei, and Michele Samorani. 2020. The evolutionary trajectories of peer-produced artifacts: Group composition, the trajectories' exploration, and the quality of artifacts. *MIS Quarterly* 44, 4 (2020), 2013–2053.
- Simran Arora, Avner May, Jian Zhang, and Christopher Ré. 2020. Contextual embeddings: When are they worth it? In *Proceedings of the 2020 Annual Meeting of the Association for Computational Linguistics. ACL*, 2650–2663.
- Sofia Bapna, Mary J. Benner, and Liangfei Qiu. 2019. Nurturing online communities: An empirical investigation. *MIS Quarterly* 43, 2 (2019), 425–452.

⁴⁰The Conference and Workshop on Neural Information Processing Systems; The International Conference on Learning Representations; The International Conference on Machine Learning; The Annual Meeting of the Association for Computational Linguistics; The Conference on Empirical Methods in Natural Language Processing.

⁴¹Content at the first three conferences would not be restricted to NLP but could involve any applications of AI and machine learning in business. For this reason, one of these three conferences would perhaps be the best place to start.

- Iz Beltagy, Kyle Lo, and Arman Cohan. 2019. SciBERT: A pretrained language model for scientific text. *Proceedings of Empirical Methods in Natural Language Processing 2019. ACL*, 3615–3620.
- Anyu Belz. 2019. DeepFake news generation: Methods, detection and wider implications. *Keynote Address at the 12th International Conference on Natural Language Generation*. Tokyo, Japan (Oct. 2019).
- Yoshua Bengio, Réjean Ducharme, Pascal Vincent, and Christian Jauvin. 2003. A neural probabilistic language model. *Journal of Machine Learning Research* 3 (2003), 1137–1155.
- Victor Benjamin, Bin Zhang, Jay F. Nunamaker Jr., and Hsinchun Chen. 2016. Examining hacker participation length in cybercriminal internet-relay-chat communities. *Journal of Management Information Systems* 33, 2 (2016), 482–510.
- Victor Benjamin, Joseph S. Valacich, and Hsinchun Chen. 2019. DICE-E: A framework for conducting darknet identification, collection, evaluation with ethics. *MIS Quarterly* 43, 1 (2019), 1–22.
- Jonathan Berant, Andrew Chou, Roy Frostig, and Percy Liang. 2013. Semantic parsing on freebase from question-answer pairs. In *Proceedings of Empirical Methods in Natural Language Processing 2013. ACL*, 1533–1544.
- Federico Bianchi, Silvia Terragni, and Dirk Hovy. 2020. Pre-training is a hot topic: Contextualized document embeddings improve topic coherence. In *Proceedings of First Workshop on Insights from Negative Results in NLP. ACL*, 32–40.
- Yonatan Bisk, Ari Holtzman, Jesse Thomason, Jacob Andreas, Yoshua Bengio, Joyce Chai, Mirella Lapata, Angeliki Lazaridou, Jonathan May, Aleksandr Nisnevich, Nicolas Pinto, and Joseph Turian. 2020. Experience grounds language. In *Proceedings of Empirical Methods in Natural Language Processing 2020. ACL*, 8718–8735.
- David M. Blei, Andrew Y. Ng, and Michael I. Jordan. 2003. Latent Dirichlet Allocation. *Journal of Machine Learning Research* 3 (2003), 993–1022.
- Ivo Blohm, Christoph Riedl, Johann Füller, and Jan Marco Leimeister. 2016. Rate or trade? Identifying winning ideas in open idea sourcing. *Information Systems Research* 27, 1 (2016), 27–48.
- Rishi Bommasani, Drew A. Hudson, Ehsan Adeli, Russ Altman, Simran Arora, Sydney von Arx, Michael S. Bernstein, Jeanette Bohg, Antoine Bosselut, Emma Brunskill, Erik Brynjolfsson, Shyamal Buch, Dallas Card, Rodrigo Castellon, Niladri Chatterji, Annie Chen, Kathleen Creel, Jared Quincy Davis, Dorottya Demszky, Chris Donahue, Moussa Doumbouya, Esin Durmus, Stefano Ermon, John Etchemendy, Kavin Ethayarajh, Li Fei-Fei, Chelsea Finn, Trevor Gale, Lauren Gillespie, Karan Goel, Noah Goodman, Shelby Grossman, Neel Guha, Tatsunori Hashimoto, Peter Henderson, John Hewitt, Daniel E. Ho, Jenny Hong, Kyle Hsu, Jing Huang, Thomas Icard, Saahil Jain, Dan Jurafsky, Pratyusha Kalluri, Siddharth Karamcheti, Geoff Keeling, Fereshte Khani, Omar Khattab, Pang Wei Koh, Mark Krass, Ranjay Krishna, Rohith Kuditipudi, Ananya Kumar, Faisal Ladhak, Mina Lee, Tony Lee, Jure Leskovec, Isabelle Levent, Xiang Lisa Li, Xuechen Li, Tengyu Ma, Ali Malik, Christopher D. Manning, Suvir Mirchandani, Eric Mitchell, Zanele Munyikwa, Suraj Nair, Avanika Narayan, Deepak Narayanan, Ben Newman, Allen Nie, Juan Carlos Nieves, Hamed Nilforoshan, Julian Nyarko, Giray Ogut, Laurel Orr, Isabel Papadimitriou, Joon Sung Park, Chris Piech, Eva Portelance, Christopher Potts, Aditi Raghunathan, Rob Reich, Hongyu Ren, Frieda Rong, Yusuf Roohani, Camilo Ruiz, Jack Ryan, Christopher Ré, Dorsa Sadigh, Shiori Sagawa, Keshav Santhanam, Andy Shih, Krishnan Srinivasan, Alex Tamkin, Rohan Taori, Armin W. Thomas, Florian Tramèr, Rose E. Wang, William Wang, Bohan Wu, Jiajun Wu, Yuhuai Wu, Sang Michael Xie, Michihiro Yasunaga, Jiaxuan You, Matei Zaharia, Michael Zhang, Tianyi Zhang, Xikun Zhang, Yuhui Zhang, Lucia Zheng, Kaitlyn Zhou, and Percy Liang. 2021. On the Opportunities and Risks of Foundation Models. arXiv:2108.07258. Retrieved from <https://arxiv.org/abs/2108.07258>.
- Susan E. Brennan. 1991. Conversation with and through computers. *User Modeling and User-Adapted Interaction* 1, 1 (1991), 67–86.
- Tom B. Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel M. Ziegler, Jeffrey Wu, Clemens Winter, Christopher Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. 2020. Language models are few-shot learners. In *Advances in Neural Information Processing Systems* 33 (*NeurIPS 2020*).
- Erik Brynjolfsson, Seth Benzell, and Daniel Rock. 2020. *Understanding and Addressing the Modern Productivity Paradox*. Research Brief, MIT Work of the Future Task Force. Massachusetts Institute of Technology, Cambridge, MA.
- Yue Cao, Hui Liu, and Xiaojun Wan. 2020. Jointly learning to align and summarize for neural cross-lingual summarization. In *Proceedings of the 2020 Annual Meeting of the Association of Computational Linguistics. ACL*, 6220–6231.
- Ankush Chatterjee, Kedhar Nath Narahari, Meghana Joshi, and Puneet Agrawal. 2019. Semeval-2019 task 3: EmoContext contextual emotion detection in text. *Proceedings of 13th International Workshop on Semantic Evaluation. ACL*, 39–48.
- Ankush Chatterjee, Umang Gupta, Manoj Kumar Chinnakotla, Radhakrishnan Srikanth, Michel Galley, and Puneet Agrawal. 2019. Understanding emotions in text using deep learning and big data. *Computers in Human Behavior* 93 (2019), 309–317.
- Michael Chau, Tim MH Li, Paul W. C. Wong, Jennifer J. Xu, Paul S. F. Yip, and Hsinchun Chen. 2020. Finding people with emotional distress in online social media: A design combining machine learning and rule-based classification. *MIS Quarterly* 44, 2 (2020), 933–953.

- Yatin Chaudhary, Pankaj Gupta, Khushbu Saxena, Vivek Kulkarni, Thomas Runkler, and Hinrich Schütze. 2020. TopicBERT for energy efficient document classification. In *Proceedings of Empirical Methods in Natural Language Processing 2020*. ACL, 1682–1690.
- Brian Cheang, Bailey Wei, David Kogan, Howey Qiu, and Masud Ahmed. 2020. Language representation models for fine-grained sentiment classification. arXiv:2005.13619. Retrieved from <https://arxiv.org/abs/2005.13619>.
- Hsinchun Chen, Roger H. L. Chiang, and Veda C. Storey. 2012. Business intelligence and analytics: From big data to big impact. *MIS Quarterly* 36, 4 (2012), 1165–1188.
- Kun Chen, Xin Li, Peng Luo, and J. Leon Zhao. 2021. News-induced dynamic networks for market signaling: Understanding impact of news on firm equity value. *Information Systems Research* 32, 2 (2021), 356–377.
- Langtao Chen, Aaron Baird, and Detmar Straub. 2019. Fostering participant health knowledge and attitudes: An econometric study of a chronic disease-focused online health community. *Journal of Management Information Systems* 36, 1 (2019), 194–229.
- Wei Chen, Bin Gu, Qiang Ye, and Kevin Xiaoguo Zhu. 2019. Measuring and managing the externality of managerial responses to online customer reviews. *Information Systems Research* 30, 1 (2019), 81–96.
- Yangbi Chen, Yun Ma, Xudong Mao, and Qing Li. 2019. Multi-task learning for abstractive and extractive summarization. *Data Science and Engineering* 4, 1 (2019), 14–23.
- Zhiyu Chen, Mohamed Trabelsi, Jeff Heflin, Yinan Xu, and Brian D. Davison. 2020. Table search using a deep contextualized language model. In *Proceedings of 43rd International ACM SIGIR Conference on Research and Development in Information Retrieval*. ACM, 589–598.
- Hichang Cho, Bart Knijnenburg, Alfred Kobsa, and Yao Li. 2018. Collective privacy management in social media: A cross-cultural validation. *ACM Transactions on Computer-Human Interaction* 25, 3 (2018), 1–33.
- Krzysztof Choromanski, Valerii Likhoshesterov, David Dohan, Xingyou Song, Andreea Gane, Tamas Sarlos, Peter Hawkins, Jared Davis, Afroz Mohiuddin, Lukasz Kaiser, David Belanger, Lucy Colwell, and Adrian Weller. 2021. Rethinking attention with performers. In *Proceedings of the International Conference on Learning Representations 2021*.
- Junyoung Chung, Caglar Gulcehre, KyungHyun Cho, and Yoshua Bengio. 2014. Empirical evaluation of gated recurrent neural networks on sequence modeling. In *Proceedings of NIPS 2014 Workshop on Deep Learning*.
- Sunghun Chung, Animesh Animesh, Kunsoo Han, and Alain Pinsonneault. 2020. Financial returns to firms’ communication actions on firm-initiated social media: Evidence from Facebook business pages. *Information Systems Research* 31, 1 (2020), 258–285.
- Peter Clark, Oren Etzioni, Daniel Khoshabi, Tushar Khot, Bhavana Dalvi Mishra, Kyle Richardson, Ashish Sabharwal, Carissa Schoenick, Oyvind Tafjord, Niket Tandon, Sumithra Bhakthavatsalam, Dirk Groeneveld, Michal Guerquin, and Michael Schmitz. 2019. From ‘F’ to ‘A’ on the NY Regents science exams: An overview of the Aristo Project. arXiv:1909.01958. Retrieved from <https://arxiv.org/abs/1909.01958>
- Ronan Collobert, Jason Weston, Léon Bottou, Michael Karlen, Koray Kavukcuoglu, and Pavel Kuksa. 2011. Natural language processing (almost) from scratch. *Journal of Machine Learning Research* 12 (2011), 2493–2537.
- Ronan Collobert and Jason Weston. 2008. A unified architecture for natural language processing: Deep neural networks with multitask learning. In *Proceedings of the International Conference on Machine Learning 2008*, 160–167.
- Alexis Conneau, Kartikay Khandelwal, Naman Goyal, Vishrav Chaudhary, Guillaume Wenzek, Francisco Guzmán, Edouard Grave, Myle Ott, Luke Zettlemoyer, and Veselin Stoyanov. 2019. Unsupervised cross-lingual representation learning at scale. In *Proceedings of Annual Meeting of the Association of Computational Linguistics 2020*. ACL, 8440–8451.
- Kristle Cortés, Ran Duchin, and Denis Sosyura. 2016. Clouded judgment: The role of sentiment in credit origination. *Journal of Financial Economics* 121, 2 (2016), 392–413.
- Zihang Dai, Zhilin Yang, Yiming Yang, Jaime Carbonell, Quoc V. Le, and Ruslan Salakhutdinov. 2019. Transformer-XL: Attentive language models beyond a fixed-length context. In *Proceedings of Annual Meeting of the Association of Computational Linguistics 2019*. ACL, 2978–2988.
- Shuyuan Deng, Zhijian James Huang, Atish P. Sinha, and Huimin Zhao. 2018. The interaction between microblog sentiment and stock return: An empirical examination. *MIS Quarterly* 42, 3 (2018), 895–918.
- Harm de Vries, Dzmitry Bahdanau, and Christopher Manning. 2020. Towards ecologically valid research on language user interfaces. arXiv:2007.14435. Retrieved from <https://arxiv.org/abs/2007.14435>.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. BERT: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of the Annual Meeting of the North American Chapter of the Association of Computational Linguistics 2019*. ACL, 4171–4186.
- Wei Dong, Shaoyi Liao, and Zhongju Zhang. 2018. Leveraging financial social media data for corporate fraud detection. *Journal of Management Information Systems* 35, 2 (2018), 461–487.
- Xiangyu Duan, Hoongfei Yu, Mingming Yin, Min Zhang, Weihua Luo, and Yue Zhang. 2019. Contrastive attention mechanism for abstractive sentence summarization. In *Proceedings of Empirical Methods in Natural Language Processing 2019*. ACL, 3044–3053.

- Mohammadreza Ebrahimi, Yidong Chai, Sagar Samtani, and Hsinchun Chen. 2021. Cross-lingual cybersecurity analytics in the international dark web with adversarial deep representation learning. *MIS Quarterly* forthcoming.
- Tim Fountaine, Brian McCarthy, and Tamim Saleh. 2019. Building the AI-powered organization. *Harvard Business Review* 97, 4 (2019), 62–73.
- Joey F. George, Manjul Gupta, Gabriel Giordano, Annette M. Mills, Vanesa M. Tennant, and Carmen C. Lewis. 2018. The effects of communication media and culture on deception detection accuracy. *MIS Quarterly* 42, 2 (2018), 551–575.
- Manoochehr Ghiassi, David Zimbra, and Sean Lee. 2016. Targeted Twitter sentiment analysis for brands using supervised feature engineering and the dynamic architecture for artificial neural networks. *Journal of Management Information Systems* 33, 4 (2016), 1034–1058.
- Deepanway Ghosal, Navonil Majumder, Alexander Gelbukh, Rada Mihalcea, and Soujanya Poria. 2020. COSMIC: COmmonSense knowledge for eMotion identification in conversations. arXiv:2010.02795. Retrieved from <https://arxiv.org/abs/2010.02795>.
- Jing Gong, Vibhanshu Abhishek, and Beibei Li. 2017. Examining the impact of keyword ambiguity on search advertising performance: A topic model approach. *MIS Quarterly* 43, 3 (2017), 805–829.
- Graciela H. Gonzalez, Tasnia Tahsin, Britton C. Goodale, Anna C. Greene, and Casey S. Greene. 2016. Recent advances and emerging applications in text and data mining for biomedical discovery. *Briefings in Bioinformatics* 17, 1 (2016), 33–42.
- Ross Gruetzemacher, David Paradice, and Kang Bok Lee. 2020. Forecasting extreme labor displacement: A survey of AI practitioners. *Technological Forecasting and Social Change*, 161 (2020). DOI : <https://doi.org/10.1016/j.techfore.2020.120323>
- Yang Gu and Yanke Hu. 2019. Extractive summarization with very deep pretrained language model. *International Journal of Artificial Intelligence and Applications* 10, 2 (2019), 27–32.
- Bikash Gyawali, Lucas Anastasiou, and Petr Knuth. 2020. Deduplication of scholarly documents using locality sensitive hashing and word embeddings. In *Proceedings of 12th Language Resources and Evaluation Conference*. ACL, 901–910.
- Jie Hao, Xing Wang, Shuming Shi, Jinfeng Zhang, and Zhaopeng Tu. 2019. Multi-granularity self-attention for neural machine translation. In *Proceedings of Empirical Methods in Natural Language Processing 2019*. ACL, 887–897.
- Pengcheng He, Xiaodong Liu, Jianfeng Gao, and Weizhu Chen. 2021. DeBERTa: Decoding-enhanced BERT with disentangled attention. In *Proceedings of the International Conference on Learning Representations 2021*.
- Irina Heimbach and Oliver Hinz. 2018. The impact of sharing mechanism design on content sharing in online social networks. *Information Systems Research* 29, 3 (2018), 592–611.
- Sepp Hochreiter and Jürgen Schmidhuber. 1997. Long short-term memory. *Neural Computation* 9, 8 (1997), 1735–1780.
- Neil Houlsby, Andrei Giurgiu, Stanislaw Jastrzebski, Bruna Morrone, Quentin De Laroussilhe, Andrea Gesmundo, Mona Attariyan, and Sylvain Gelly. 2019. Parameter-efficient transfer learning for NLP. In *Proceedings of the International Conference on Machine Learning 2019*. 2790–2799.
- Jeremy Howard and Sebastian Ruder. 2018. Universal language model fine-tuning for text classification. arXiv:1801.06146. Retrieved from <https://arxiv.org/abs/1801.06146>.
- Alexander Hoyle, Pranav Goel, and Philip Resnik. 2020. Improving neural topic models using knowledge distillation. In *Proceedings of Empirical Methods in Natural Language Processing 2020*. ACL, 1752–1771.
- Jiaxin Huang, Yu Meng, Fang Guo, Heng Ji, and Jiawei Han. 2020. Weakly-supervised aspect-based sentiment analysis via joint aspect-sentiment topic embedding. arXiv:2010.06705. <https://arxiv.org/abs/2010.06705>.
- Jianxiong Huang, Wai Fong Boh, and Kim Huat Goh. 2017. A temporal study of the effects of online opinions: Information sources matter. *Journal of Management Information Systems* 34, 4 (2017), 1169–1202.
- Kuang-Yuan Huang, Indushobha Chengalur-Smith, and Alain Pinsonneault. 2019. Sharing is caring: Social support provision and companionship activities in healthcare virtual support communities. *MIS Quarterly* 43, 2 (2019), 395–424.
- Ni Huang, Yili Hong, and Gordon Burtch. 2017. Social network integration and user content generation: Evidence from natural experiments. *MIS Quarterly* 41, 4 (2017), 1035–1058.
- Po-Sen Huang, Huan Zhang, Ray Jiang, Robert Stanforth, Johannes Welbl, Jack Rae, Vishal Maini, Dani Yogatama, and Pushmeet Kohli. 2019. Reducing sentiment bias in language models via counterfactual evaluation. In *Proceedings of Empirical Methods in Natural Language Processing 2019*. ACL, 65–83.
- Elina H. Hwang, Param Vir Singh, and Linda Argote. 2019. Jack of all, master of some: Information network and innovation in crowdsourcing communities. *Information Systems Research* 30, 2 (2019), 389–410.
- Fuwei Jiang, Joshua Lee, Xiumin Martin, and Guofu Zhou. 2019. Manager sentiment and stock returns. *Journal of Financial Economics* 132, 1 (2019), 126–149.
- Hanqi Jin, Tianming Wang, and Xiaojun Wan. 2020. Multi-granularity interaction network for extractive and abstractive multi-document summarization. In *Proceedings of Annual Meeting of the Association of Computational Linguistics 2020*. ACL, 6244–6254.
- Jared Kaplan, Sam McCandlish, Tom Henighan, Tom B. Brown, Benjamin Chess, Rewon Child, Scott Gray, Alec Radford, Jeffrey Wu, and Dario Amodei. 2020. Scaling laws for neural language models. arXiv:2001.08361. Retrieved from <https://arxiv.org/abs/2001.08361>.

- Pei Ke, Haozhe Ji, Siyang Liu, Xiaoyan Zhu, and Minlie Huang. 2020. SentiLARE: Linguistic knowledge enhanced language representation for sentiment analysis. In *Proceedings of Empirical Methods in Natural Language Processing 2020*. ACL, 6975–6988.
- Nitish Shirish Keskar, Bryan McCann, Lav R. Varshney, Caiming Xiong, and Richard Socher. 2019. CTRL: A conditional transformer language model for controllable generation. arXiv:1909.05858. Retrieved from <https://arxiv.org/abs/1909.05858>.
- Daniel Khashabi, Sewon Min, Tushar Khot, Ashish Sabharwal, Oyvind Tafjord, Peter Clark, and Hannaneh Hajishirzi. 2020. UnifiedQA: Crossing format boundaries with a single QA system. In *Proceedings of Empirical Methods in Natural Language Processing 2020*. ACL, 1896–1907.
- Warut Khern-am-nuai, Karthik Kannan, and Hossein Ghasemkhani. 2018. Extrinsic versus intrinsic rewards for contributing reviews in an online platform. *Information Systems Research* 29, 4 (2018), 871–892.
- Nikita Kitaev, Łukasz Kaiser, and Anselm Levskaya. 2020. Reformer: The efficient transformer. In *Proceedings of the International Conference on Learning Representations 2020*.
- Mathias Kraus, and Stefan Feuerriegel. 2017. Decision support from financial disclosures with deep neural networks and transfer learning. *Decision Support Systems* 104 (2017), 38–48.
- Mathias Kraus, Stefan Feuerriegel, and Asil Oztekin. 2020. Deep learning in business analytics and operations research: Models, applications and managerial implications. *European Journal of Operational Research* 281, 3 (2020), 628–641.
- Wojciech Kryściński, Nitish Shirish Keskar, Bryan McCann, Caiming Xiong, and Richard Socher. 2019. Neural text summarization: A critical evaluation. In *Proceedings of Empirical Methods in Natural Language Processing 2019*. ACL, 540–551.
- Naveen Kumar, Deepak Venugopal, Liangfei Qiu, and Subodha Kumar. 2019. Detecting anomalous online reviewers: An unsupervised approach using mixture models. *Journal of Management Information Systems* 36, 4 (2019), 1313–1346.
- Theodoros Lappas, Gaurav Sabnis, and Georgios Valkanias. 2016. The impact of fake reviews on online visibility: A vulnerability assessment of the hotel industry. *Information Systems Research* 27, 4 (2016), 940–961.
- Zhenzhong Lan, Mingda Chen, Sebastian Goodman, Kevin Gimpel, Piyush Sharma, and Radu Soricut. 2019. ALBERT: A lite BERT for self-supervised learning of language representations. In *Proceedings of the International Conference on Learning Representations 2019*.
- Yann LeCun, Yoshua Bengio, and Geoffrey Hinton. 2015. Deep learning. *Nature* 521 (2015), 436–444.
- Gene Moo Lee, Shu He, Joowon Lee, and Andrew B. Whinston. 2020. Matching mobile applications for cross-promotion. *Information Systems Research* 31, 3 (2020), 865–891.
- Mike Lewis, Yinhan Liu, Naman Goyal, Marjan Ghazvininejad, Abdelrahman Mohamed, Omer Levy, Ves Stoyanov, and Luke Zettlemoyer. 2020a. BART: Denoising sequence-to-sequence pre-training for natural language generation, translation, and comprehension. In *Proceedings of the Annual Meeting of the Association of Computational Linguistics 2020*.
- Mike Lewis, Marjan Ghazvininejad, Gargi Ghosh, Armen Aghajanyan, Sida Wang, and Luke Zettlemoyer. 2020b. Pre-training via paraphrasing. In *Advances in Neural Information Processing Systems 33 (NeurIPS 2020)*.
- Patrick Lewis, Ethan Perez, Aleksandra Piktus, Fabio Petroni, Vladimir Karpukhin, Naman Goyal, Heinrich Küttler, Mike Lewis, Wen-tau Yih, Tim Rocktäschel, Sebastian Riedel, and Douwe Kiela. 2020c. Retrieval-augmented generation for knowledge-intensive NLP tasks. In *Advances in Neural Information Processing Systems 33 (NeurIPS 2020)*.
- Jingjing Li, Kai Larsen, and Ahmed Abbasi. 2020. TheoryOn: A design framework and system for unlocking behavioral knowledge through ontology learning. *MIS Quarterly* 44, 4 (2020), 1733–1772.
- Menggang Li, Wenrui Li, Fang Wang, Xiaojun Jia, and Guangwei Rui. 2021. Applying BERT to analyze investor sentiment in stock market. *Neural Computing and Applications* 33, 10 (2021), 4663–4676.
- Weifeng Li, Hsinchun Chen, and Jay F. Nunamaker Jr. 2016. Identifying and profiling key sellers in cyber carding community: AZSecure text mining system. *Journal of Management Information Systems* 33, 4 (2016), 1059–1086.
- Xin Li, Lidong Bing, Wenxuan Zhang, and Wai Lam. 2019. Exploiting BERT for end-to-end aspect-based sentiment analysis. In *Proceedings of 2019 EMNLP Workshop W-NUT*. ACL, 34–41.
- Huigang Liang, Yajiong Xue, Alain Pinsonneault, and Yu Wu. 2019. What users do besides problem-focused coping when facing IT security threats: An emotion-focused coping perspective. *MIS Quarterly* 43, 2 (2019), 373–394.
- Yaobo Liang, Nan Duan, Yeyun Gong, Ning Wu, Fenfei Guo, Weizhen Qi, Ming Gong, Linjun Shou, Daxin Jiang, Guihong Cao, Xiaodong Fan, Ruofei Zhang, Rahul Agrawal, Edward Cui, Sining Wei, Taroon Bharti, Ying Qiao, Jiun-Hung Chen, Winnie Wu, Shuguang Liu, Fan Yang, Daniel Campos, Rangan Majumder, and Ming Zho. 2020. XGLUE: A new benchmark dataset for cross-lingual pre-training, understanding and generation. In *Proceedings of Empirical Methods in Natural Language Processing 2020*. ACL, 6008–6018.
- Opher Lieber, Or Sharir, Barak Lenz, and Yoav Shoham. 2021. Jurassic-1: Technical details and evaluation. *White Paper, AI21 Labs*. Tel Aviv, Israel.
- Xiao Liu, Bin Zhang, Anjana Susarla, and Rema Padman. 2020a. Go to YouTube and call me in the morning: Use of social media for chronic conditions. *MIS Quarterly* 44, 1 (2020), 257–283.

- Xiaodong Liu, Pengcheng He, Weizhu Chen, and Jianfeng Gao. 2019a. Multi-task deep neural networks for natural language understanding. In *Proceedings of the Annual Meeting of the Association of Computational Linguistics 2019. ACL*, 4487–4496.
- Xiaomo Liu, G. Alan Wang, Weiguo Fan, and Zhongju Zhang. 2020b. Finding useful solutions in online knowledge communities: A theory-driven design and multilevel analysis. *Information Systems Research* 31, 3 (2020), 731–752.
- Yang Liu and Mirella Lapata. 2019b. Text summarization with pretrained encoders. In *Proceedings of Empirical Methods in Natural Language Processing 2019. ACL*, 3730–3740.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019c. RoBERTa: A robustly optimized BERT pretraining approach. *arXiv:1907.11692*. Retrieved from <https://arxiv.org/abs/1907.11692>.
- Yuanyang Liu, Gautam Pant, and Olivia R. L. Sheng. 2020c. Predicting labor market competition: Leveraging interfirm network and employee skills. *Information Systems Research* 31, 4 (2020), 1443–1466.
- Yao Lu, Yue Dong, and Laurent Charlin. 2020. Multi-XScience: A large-scale dataset for extreme multi-document summarization of scientific articles. In *Proceedings of Empirical Methods in Natural Language Processing 2020. ACL*, 8068–8074.
- Xinyao Ma, Maarten Sap, Hannah Rashkin, and Yejin Choi. 2020. PowerTransformer: Unsupervised controllable revision for biased language correction. In *Proceedings of Empirical Methods in Natural Language Processing 2020. ACL*, 7426–7441.
- Feng Mai, Zhe Shan, Qing Bai, Xin Wang, and Roger H. L. Chiang. 2018. How does social media impact Bitcoin value? A test of the silent majority hypothesis. *Journal of Management Information Systems* 35, 1 (2018), 19–52.
- Christopher Manning and Hinrich Schütze. 1999. *Foundations of Statistical Natural Language Processing*. MIT Press, Cambridge, MA.
- Joshua Maynez, Shashi Narayan, Bernd Bohnet, and Ryan McDonald. 2020. On faithfulness and factuality in abstractive summarization. In *Proceedings of Empirical Methods in Natural Language Processing 2020. ACL*, 1906–1919.
- Bryan McCann, James Bradbury, Caiming Xiong, and Richard Socher. 2017. Learned in translation: Contextualized word vectors. In *Advances in Neural Information Processing Systems 30 (NIPS 2017)*.
- Jorge Mejia, Shawn Mankad, and Anandasivam Gopal. 2019. A for effort? Using the crowd to identify moral hazard in New York City restaurant hygiene inspections. *Information Systems Research* 30, 4 (2019), 1363–1386.
- Tomas Mikolov, Ilya Sutskever, Kai Chen, Greg S. Corrado, and Jeff Dean. 2013a. Distributed representations of words and phrases and their compositionality. In *Advances in Neural Information Processing Systems 26 (NIPS 2013)*. 3111–3119.
- Tomas Mikolov, Kai Chen, Greg Corrado, and Jeffrey Dean. 2013b. Efficient estimation of word representations in vector space. In *Proceedings of the International Conference on Learning Representations 2013*.
- Tomas Mikolov, Armand Joulin, Sumit Chopra, Michael Mathieu, and Marc’Aurelio Ranzato. 2014. Learning longer memory in recurrent neural networks. *arXiv:1412.7753*. Retrieved from <https://arxiv.org/abs/1412.7753>.
- Antonio Moreno, and Teófilo Redondo. 2016. Text analytics: The convergence of big data and artificial intelligence. *International Journal of Interactive Multimedia and Artificial Intelligence* 3, 6 (2016), 57–64.
- Morgan R. Frank, David Autor, James E. Bessen, Erik Brynjolfsson, Manuel Cebrian, David J. Deming, Maryann Feldman, Matthew Groh, Jos ‘e Lobo, Esteban Moro, Dashun Wang, Hyejin Youn, and Iyad Rahwan. 2019. Toward understanding the impact of artificial intelligence on labor. In *Proceedings of the National Academy of Sciences* 116, 14 (2019), 6531–6539.
- Reza Mousavi and Bin Gu. 2019. The impact of Twitter adoption on lawmakers’ voting orientations. *Information Systems Research* 30, 1 (2019), 133–153.
- Reza Mousavi, Monica Johar, and Vijay S. Mookerjee. 2020. The voice of the customer: Managing customer care in Twitter. *Information Systems Research* 31, 2 (2020), 340–360.
- Martin Müller, Marcel Salathé, and Per E. Kummervold. 2020. COVID-Twitter-BERT: A natural language processing model to analyse COVID-19 content on Twitter. *arXiv:2005.07503*. Retrieved from <https://arxiv.org/abs/2005.07503>.
- Kevin P. Murphy. 2012. *Machine Learning: A Probabilistic Perspective*. MIT Press, Cambridge, MA.
- Shashi Narayan, Shay B. Cohen, and Mirella Lapata. 2019. What is this article about? Extreme summarization with topic-aware convolutional neural networks. *Journal of Artificial Intelligence Research*, 66, 243–278.
- Usman Naseem, Imran Razzak, Katarzyna Musiał, and Muhammad Imran. 2020. Transformer based deep intelligent contextual embedding for Twitter sentiment analysis. *Future Generation Computer Systems* 113 (2020), 58–69.
- Eric W. T. Ngai and Philip Tin Yun Lee. 2016. A review of the literature on applications of text mining in policy making. In *Proceedings of the Pacific Asia Conference on Information Systems. AIS*.
- Dat Quoc Nguyen, Thanh Vu, and Anh Tuan Nguyen. 2020. BERTweet: A pre-trained language model for English Tweets. In *Proceedings of Empirical Methods in Natural Language Processing 2020: Demonstrations. ACL*, 9–14.
- Yang Pan, Peng Huang, and Anandasivam Gopal. 2019. Storm clouds on the horizon? New entry threats and R&D investments in the US IT industry. *Information Systems Research* 30, 2 (2019), 540–562.
- Jinuk Park, Chanhee Park, Jeongwoo Kim, Minsoo Cho, and Sanghyun Park. 2019. ADC: Advanced document clustering using contextualized representations. *Expert Systems with Applications* 137 (2019), 157–166.
- Jeffrey Pennington, Richard Socher, and Christopher D. Manning. 2014. GloVe: Global vectors for word representation. In *Proceedings of Empirical Methods in Natural Language Processing. ACL*, 1532–1543.

- Matthew E. Peters, Mark Neumann, Mohit Iyyer, Matt Gardner, Christopher Clark, Kenton Lee, and Luke Zettlemoyer. 2018. Deep contextualized word representations. In *Proceedings of the Annual Meeting of the North American Chapter of the Association of Computational Linguistics* 2018. ACL, 2227–2237.
- Alec Radford, Karthik Narasimhan, Tim Salimans, and Ilya Sutskever. 2018. Improving language understanding by generative pre-training. (11, 2018). Retrieved Dec. 11, 2021 from https://s3-us-west-2.amazonaws.com/openai-assets/research-covers/language-unsupervised/language_understanding_paper.pdf.
- Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. 2019. Language models are unsupervised multitask learners. (14, 2019). Retrieved Dec. 11, 2021 from https://d4mucfpxsywv.cloudfront.net/better-language-models/language_models_are_unsupervised_multitask_learners.pdf.
- Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J. Liu. 2020. Exploring the limits of transfer learning with a unified text-to-text transformer. *Journal of Machine Learning Research* 21 (2020), 1–67.
- Nils Reimers, Benjamin Schiller, Tilman Beck, Johannes Daxenberger, Christian Stab, and Iryna Gurevych. 2019. Classification and clustering of arguments with contextualized word embeddings. In *Proceedings of the Annual Meeting of the Association of Computational Linguistics* 2019. ACL, 567–578.
- Lauren Rhue and Arun Sundararajan. 2019. Playing to the crowd? Digital visibility and the social dynamics of purchase disclosure. *MIS Quarterly* 43, 4 (2019), 1127–1141.
- Matthew D. Rocklage and Russell H. Fazio. 2020. The enhancing versus backfiring effects of positive emotion in consumer reviews. *Journal of Marketing Research* 57, 2 (2020), 332–352.
- Stephen Roller, Emily Dinan, Naman Goyal, Da Ju, Mary Williamson, Yinhan Liu, Jing Xu, Myle Ott, Kurt Shuster, Eric M. Smith, Y-Lan Boureau, and Jason Weston. 2020a. Recipes for building an open-domain chatbot. In *Proceedings of the 2020 Conference of the European Chapter of the Association for Computational Linguistics*. ACL, 300–325.
- Stephen Roller, Y-Lan Boureau, Jason Weston, Antoine Bordes, Emily Dinan, Angela Fan, David Gunning, Da Ju, Margaret Li, Spencer Poff, Pratik Ringshia, Kurt Shuster, Eric Michael Smith, Arthur Szlam, Jack Urbanek, and Mary Williamson. 2020b. Open-domain conversational agents: Current progress, open problems, and future directions. *arXiv:2006.12442*. Retrieved from <https://arxiv.org/abs/2006.12442>.
- Danish H. Saifee, Indranil R. Bardhan, Atanu Lahiri, and Zhiqiang Zheng. 2019. Adherence to clinical guidelines, electronic health record use, and online reviews. *Journal of Management Information Systems* 36, 4 (2019), 1071–1104.
- Shimi Salant and Jonathan Berant. 2018. Contextualized word representations for reading comprehension. In *Proceedings of the Annual Meeting of the North American Chapter of the Association of Computational Linguistics*. ACL, 554–559.
- Markus Salo, Markus Mykkänen, and Riitta Hekkala. 2020. The interplay of IT users’ coping strategies: Uncovering momentary emotional load, routes, and sequences. *MIS Quarterly* 44, 3 (2020), 1143–1175.
- Sagar Samtani, Ryan Chinn, Hsinchun Chen, and Jay F. Nunamaker Jr. 2017. Exploring emerging hacker assets and key hackers for proactive cyber threat intelligence. *Journal of Management Information Systems* 34, 4 (2017), 1023–1053.
- Sagar Samtani, Hongyi Zhu, Balaji Padmanabhan, Yidong Chai, and Hsinchun Chen. 2020. Deep learning for information systems research. *arXiv:2010.05774*. Retrieved from <https://arxiv.org/abs/2010.05774>.
- Sagar Samtani, Yidong Chai, and Hsinchun Chen. 2021. Linking exploits from the dark web to known vulnerabilities for proactive cyber threat intelligence: An attention-based deep structured semantic model. *MIS Quarterly* forthcoming.
- Victor Sanh, Lysandre Debut, Julien Chaumond, and Thomas Wolf. 2019. DistilBERT, a distilled version of BERT: Smaller, faster, cheaper and lighter. In *Proceedings of the 2019 NeurIPS Workshop on Energy Efficient Machine Learning and Cognitive Computing*.
- Timo Schick and Hinrich Schütze. 2020. It’s not just size that matters: Small language models are also few-shot learners. *arXiv:2009.07118*. Retrieved from <https://arxiv.org/abs/2009.07118>.
- Dan Schwartz, Mariya Toneva, and Leila Wehbe. 2019. Inducing brain-relevant bias in natural language processing models. In *Advances in Neural Information Processing Systems* 32 (NeurIPS 2019).
- Thomas Scialom, Paul-Alexis Dray, Sylvain Lamprier, Benjamin Piwowarski, and Jacopo Staiano. 2020. MLSUM: The multi-lingual summarization corpus. In *Proceedings of Empirical Methods in Natural Language Processing* 2020. ACL, 8051–8067.
- Lütfi Kerem Şenel, İhsan Utlü, Veyisel Yücesoy, Aykut Koc, and Tolga Cukur. 2018. Semantic structure and interpretability of word embeddings. *IEEE/ACM Transactions on Audio, Speech, and Language Processing* 26, 10 (2018), 1769–1779.
- Donghui Shi, Jian Guan, Jozef Zurada, and Andrew Manikas. 2017. A data-mining approach to identification of risk factors in safety management systems. *Journal of Management Information Systems* 34, 4 (2017), 1054–1081.
- Zhan Shi, Gene Moo Lee, and Andrew B. Whinston. 2016. Toward a better measure of business proximity: Topic modeling for industry intelligence. *MIS Quarterly* 40, 4 (2016), 1035–1056.
- Donghyuk Shin, Shu He, Gene Moo Lee, Andrew B. Whinston, Suleyman Cetintas, and Kuang-Chih Lee. 2020. Enhancing social media analysis with visual data analytics: A deep learning approach. *MIS Quarterly* 44, 4 (2020), 1459–1492.
- Suzanna Sia, Ayush Dalmia, and Sabrina J. Mielke. 2020. Tired of topic models? *Clusters of pretrained word embeddings make for fast and good topics too!* In *Proceedings of Empirical Methods in Natural Language Processing* 2020. ACL, 1728–1736.

- Aditya Siddhant, Junjie Hu, Melvin Johnson, Orhan Firat, and Sebastian Ruder. 2020. Xtreme: A massively multilingual multi-task benchmark for evaluating cross-lingual generalization. In *Proceedings of the International Conference on Machine Learning 2020*. 4411–4421.
- Michael Siering, Jascha-Alexander Koch, and Amit V. Deokar. 2016. Detecting fraudulent behavior on crowdfunding platforms: The role of linguistic and content-based cues in static and dynamic contexts. *Journal of Management Information Systems* 33, 2 (2016), 421–455.
- David Silver, Julian Schrittwieser, Karen Simonyan, Ioannis Antonoglou, Aja Huang, Arthur Guez, Thomas Hubert, Lucas Baker, Matthew Lai, Adrian Bolton, Yutian Chen, Timothy Lillicrap, Fan Hui, Laurent Sifre, George van den Driessche, Thore Graepel, and Demis Hassabis. 2017. Mastering the game of Go without human knowledge. *Nature* 550 (2017), 354–359.
- Noah A. Smith. 2020. Contextual word representations: Putting words into computers. *Communications of the ACM* 63, 6 (2020), 66–74.
- Tingting Song, Jinghua Huang, Yong Tan, and Yifan Yu. 2019. Using user- and marketer-generated content for box office revenue prediction: Differences between microblogging and third-party platforms. *Information Systems Research* 30, 1 (2019), 191–203.
- Nisan Stiennon, Long Ouyang, Jeff Wu, Daniel M. Ziegler, Ryan Lowe, Chelsea Voss, Alec Radford, Dario Amodei, and Paul Christiano. 2020. Learning to summarize with human feedback. In *Advances in Neural Information Processing Systems* 33 (NeurIPS 2020).
- Chi Sun, Xipeng Qiu, Yige Xu, and Xuanjing Huang. 2019. How to fine-tune BERT for text classification?. In *Proceedings of China National Conference on Chinese Computational Linguistics*. Springer, 194–206.
- Yu Sun, Shuohuan Wang, Shikun Feng, Siyu Din, Chao Pan, Junyuan Shan, Jiaxiang Li, Xuyi Chen, Yanbin Zhao, Yuxiang Lu, Weixin Liu, Zhihua Wu, Weibao Gong, Jianzhong Liang, Zhizhou Shang, Peng Sun, Wei Liu, Xuan Ouyang, Dianhai Yu, Hao Tian, Hua Wu, and Haifeng Wang. 2021. ERNIE 3.0: Large-scale knowledge enhanced pre-training for language understanding and generation. *arXiv:2107.02137*. Retrieved from <https://arxiv.org/abs/2107.02137>.
- Hao Tan and Mohit Bansal. 2019. LXMERT: Learning cross-modality encoder representations from transformers. In *Proceedings of Empirical Methods in Natural Language Processing*, ACL, 5100–5111.
- Hao Tan and Mohit Bansal. 2020. Vokenization: Improving language understanding with contextualized, visual-grounded supervision. In *Proceedings of Empirical Methods in Natural Language Processing*, ACL, 2066–2080.
- Yi Tay, Mostafa Dehghani, Jai Gupta, Dara Bahri, Vamsi Aribandi, Zhen Qin, and Donald Metzler. 2021. Are pre-trained convolutions better than pre-trained transformers? In *Proceedings of the Annual Meeting of the Association of Computational Linguistics 2021*.
- Alan. M. Turing. 1950. Computing machinery and intelligence. *Mind* 59, 236 (1950), 433–460.
- Wietske Van Osch and Charles W. Steinfield. 2018. Strategic visibility in enterprise social media: Implications for network formation and boundary spanning. *Journal of Management Information Systems* 35, 2 (2018), 647–682.
- 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* 30 (NIPS 2017).
- Sriker Velichety, Sudha Ram, and Jesse Bockstedt. 2019. Quality assessment of peer-produced content in knowledge repositories using development and coordination activities. *Journal of Management Information Systems* 36, 2 (2019), 478–512.
- Alex Wang, Amanpreet Singh, Julian Michael, Felix Hill, Omer Levy, and Samuel R. Bowman. 2018a. GLUE: A multi-task benchmark and analysis platform for natural language understanding. In *Proceedings of EMNLP Workshop on BlackBox NLP*, ACL, 353–355.
- Alex Wang, Yada Pruksachatkun, Nikita Nangia, Amanpreet Singh, Julian Michael, Felix Hill, Omer Levy, and Samuel R. Bowman. 2019a. SuperGLUE: A stickier benchmark for general-purpose language understanding systems. In *Advances in Neural Information Processing Systems* 32 (NeurIPS 2019).
- Quan Wang, Beibei Li, and Param Vir Singh. 2018b. Copycats vs. original mobile apps: A machine learning copycat-detection method and empirical analysis. *Information Systems Research* 29, 2 (2018), 273–291.
- Xuequn Wang and Zilong Liu. 2019b. Online engagement in social media: A cross-cultural comparison. *Computers in Human Behavior* 97 (2019), 137–150.
- Zhao Wang, Cuiqing Jiang, Huimin Zhao, and Yong Ding. 2020. Mining semantic soft factors for credit risk evaluation in peer-to-peer lending. *Journal of Management Information Systems* 37, 1 (2020), 282–308.
- Zirui Wang, Yulia Tsvetkov, Orhan Firat, and Yuan Cao. 2021a. Gradient vaccine: Investigating and improving multi-task optimization in massively multilingual models. In *Proceedings of the International Conference on Learning Representations 2021*.
- Zirui Wang, Adams Wei Yu, Orhan Firat, and Yuan Cao. 2021b. Towards zero-label language learning. *arXiv:2109.09193*. Retrieved from <https://arxiv.org/abs/2109.09193>.

- Jason Wei, Maarten Bosma, Vincent Y. Zhao, Kelvin Guu, Adams Wei Yu, Brian Lester, Nan Du, Andrew M. Dai, and Quoc V. Le. 2021. Finetuned language models are zero-shot learners. *arXiv:2109.01652*. Retrieved from <https://arxiv.org/abs/2109.01652>.
- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Remi Louf, Morgan Funtowicz, Joe Davison, Sam Shleifer, Clara Ma, Yacine Jernite, Julien Plu, Canwen Xu, Teven Le Scao, Sylvain Gugger, Mariama Drame, Quentin Lhoest, and Alexander M. Rush. 2020. Transformers: State-of-the-art natural language processing. In *Proceedings of Empirical Methods in Natural Language Processing 2020: Demonstrations*. ACL, 38–45.
- Chien-Sheng Wu, Steven Hoi, Richard Socher, and Caiming Xiong. 2020. TOD-BERT: Pre-trained natural language understanding for task-oriented dialogues. In *Proceedings of Empirical Methods in Natural Language Processing 2020*. ACL, 917–929.
- Ji Wu, Liqiang Huang, and J. Leon Zhao. 2019. Operationalizing regulatory focus in the digital Age: Evidence from an E-commerce context. *MIS Quarterly* 43, 3 (2019), 745–764.
- Jeff Wu, Long Ouyang, Daniel M. Ziegler, Nisan Stiennon, Ryan Lowe, Jan Leike, and Paul Christiano. 2021. Recursively summarizing books with human feedback. *arXiv:2109.10862*. Retrieved from <https://arxiv.org/abs/2109.10862>.
- Patrick Xia, Shijie Wu, and Benjamin Van Durme. 2020. Which* BERT? A survey organizing contextualized encoders. In *Proceedings of Empirical Methods in Natural Language Processing 2020*. ACL, 7516–7533.
- Hu Xu, Bing Liu, Lei Shu, and Philip S. Yu. 2019. BERT post-training for review reading comprehension and aspect-based sentiment analysis. In *Proceedings of the Annual Meeting of the North American Chapter of the Association of Computational Linguistics 2019*. ACL, 2324–2335.
- Kelvin Xu, Jimmy Ba, Ryan Kiros, Kyunghyun Cho, Aaron Courville, Ruslan Salakhudinov, Rich Zemel, and Yoshua Bengio. 2015. Show, attend and tell: Neural image caption generation with visual attention. In *Proceedings of the International Conference on Machine Learning*. 2048–2057.
- Yumo Xu and Mirella Lapata. 2020. Coarse-to-fine query focused multi-document summarization. In *Proceedings of Empirical Methods in Natural Language Processing 2020*. ACL, 3632–3645.
- Mochen Yang, Gediminas Adomavicius, Gordon Burtch, and Yuqing Ren. 2018. Mind the gap: Accounting for measurement error and misclassification in variables generated via data mining. *Information Systems Research* 29, 1 (2018), 4–24.
- Ruosong Yang, Jiannong Cao, Zhiyuan Wen, Youzheng Wu, and Xiaodong He. 2020. Enhancing automated essay scoring performance via cohesion measurement and combination of regression and ranking. In *Proceedings Empirical Methods in Natural Language Processing 2020: Findings*. ACL, 1560–1569.
- Zhilin Yang, Zihang Dai, Yiming Yang, Jaime Carbonell, Russ R. Salakhutdinov, and Quoc V. Le. 2019. XLNet: Generalized autoregressive pretraining for language understanding. In *Advances in Neural Information Processing Systems 32 (NurIPS 2019)*.
- Liang Yao, Zhe Jin, Chengsheng Mao, Yin Zhang, and Yuan Luo. 2019. Traditional Chinese medicine clinical records classification with BERT and domain specific corpora. *Journal of the American Medical Informatics Association* 26, 12 (2019), 1632–1636.
- Eunae Yoo, Bin Gu, and Elliot Rabinovich. 2019. Diffusion on social media platforms: A point process model for interaction among similar content. *Journal of Management Information Systems* 36, 4 (2019), 1105–1141.
- Holly Young. 2015. The digital language divide. *The Guardian*. Retrieved from <http://labs.theguardian.com/digital-language-divide/>.
- Wei T. Yue, Qiu-Hong Wang, and Kai-Lung Hui. 2019. See no evil, hear no evil? Dissecting the impact of online hacker forums. *MIS Quarterly* 43, 1 (2019), 73–95.
- Manzil Zaheer, Guru Guruganesh, Avinava Dubey, Joshua Ainslie, Chris Alberti, Santiago Ontanon, Philip Pham, Anirudh Ravula, Qifan Wang, Li Yang, and Amr Ahmed. 2020. Big Bird: Transformers for longer sequences. In *Advances in Neural Information Processing Systems 33 (NeurIPS 2020)*.
- Rowan Zellers, Yonatan Bisk, Ali Farhadi, and Yejin Choi. 2019. From recognition to cognition: Visual commonsense reasoning. In *Proceedings of CVPR 2019*. CVF, 6720–6731.
- Dongsong Zhang, Lina Zhou, Juan Luo Kehoe, and Isil Yakut Kilic. 2016. What online reviewer behaviors really matter? Effects of verbal and nonverbal behaviors on detection of fake online reviews. *Journal of Management Information Systems* 33, 2 (2016), 456–481.
- Kunpeng Zhang, Siddhartha Bhattacharyya, and Sudha Ram. 2016. Large-scale network analysis for online social brand advertising. *MIS Quarterly* 40, 4 (2016), 849–868.
- Shuailiang Zhang, Hai Zhao, Yuwei Wu, Zhuosheng Zhang, Xi Zhou, and Xiang Zhou. 2020. DCMN+: Dual co-matching network for multi-choice reading comprehension. In *Proceedings of Thirty-Fourth AAAI Conference on AI (AAAI 2020)*. AAAI, 9563–9570.
- Wenli Zhang and Sudha Ram. 2020. A comprehensive analysis of triggers and risk factors for asthma based on machine learning and large heterogeneous data sources. *MIS Quarterly* 44, 1 (2020), 305–339.

- Peixiang Zhong, Di Wang, and Chunyan Miao. 2019. Knowledge-enriched transformer for emotion detection in textual conversations. In *Proceedings of Empirical Methods in Natural Language Processing 2019*. ACL, 165–176.
- Shihao Zhou, Zhilei Qiao, Qianzhou Du, G. Alan Wang, Weiguo Fan, and Xiangbin Yan. 2018. Measuring customer agility from online reviews using big data text analytics. *Journal of Management Information Systems* 35, 2 (2018), 510–539.
- Junnan Zhu, Qian Wang, Yining Wang, Yu Zhou, Jiajun Zhang, Shaonan Wang, and Chengqing Zong. 2019. NCLS: Neural cross-lingual summarization. In *Proceedings of Empirical Methods in Natural Language Processing 2019*. ACL 3054–3064.

Received March 2021; revised October 2021; accepted December 2021