

Translational NLP: A New Paradigm and General Principles for Natural Language Processing Research

Denis Newman-Griffis¹ Jill Fain Lehman² Carolyn Rosé³ Harry Hochheiser¹

¹Department of Biomedical Informatics, University of Pittsburgh, USA

²Human-Computer Interaction Institute, Carnegie Mellon University, USA

³Language Technologies Institute, Carnegie Mellon University, USA

{dnewmangriffis, harryh}@pitt.edu, {jfl, cprose}@cs.cmu.edu

Abstract

Natural language processing (NLP) research combines the study of universal principles, through basic science, with applied science targeting specific use cases and settings. However, the process of exchange between basic NLP and applications is often assumed to emerge naturally, resulting in many innovations going unapplied and many important questions left unstudied. We describe a new paradigm of *Translational NLP*, which aims to structure and facilitate the processes by which basic and applied NLP research inform one another. Translational NLP thus presents a third research paradigm, focused on understanding the challenges posed by application needs and how these challenges can drive innovation in basic science and technology design. We show that many significant advances in NLP research have emerged from the intersection of basic principles with application needs, and present a conceptual framework outlining the stakeholders and key questions in translational research. Our framework provides a roadmap for developing Translational NLP as a dedicated research area, and identifies general translational principles to facilitate exchange between basic and applied research.

1 Introduction

Natural language processing (NLP) lies at the intersection of basic science and applied technologies. However, translating innovations in basic NLP methods to successful applications remains a difficult task in which failure points often appear late in the development process, delaying or preventing potential impact in research and industry. Application challenges range widely, from changes in data distributions (Elsahar and Gallé, 2019) to computational bottlenecks (Desai et al., 2020) and integration with domain expertise (Rahman et al., 2020). When unanticipated, such challenges can be fatal to applications of new NLP methodologies, leaving exciting innovations with minimal practical

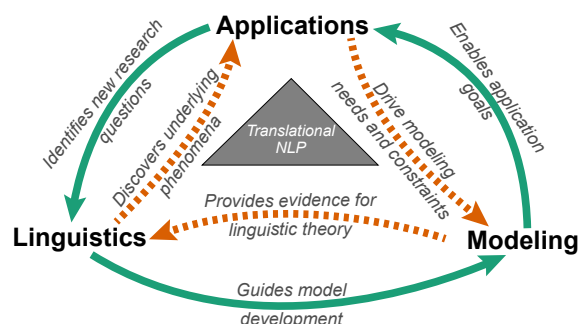


Figure 1: Interactions between linguistic theory, model development, and applications in NLP research. Solid lines indicate moving from basic research to applications, and dashed lines indicate how applied research feeds back into basic study. Translational NLP develops processes to realize this exchange.

impact. Meanwhile, real-world applications may rely on regular expressions (Anzaldi et al., 2017) or unigram frequencies (Slater et al., 2017) when more sophisticated methods would yield deeper insight. When successful translations of basic NLP insights into practical applied technologies do occur, the factors contributing to this success are rarely analyzed, limiting our ability to learn how to enable the next project and the next technology.

We argue for a third kind of NLP research, which we call *Translational NLP*. Translational NLP research aims to understand why one translation succeeds while another fails, and to develop general, reusable processes to facilitate more (and easier) translation between basic NLP advances and real-world application settings. Much NLP research already includes translational insights, but often considers them properties of a specific application rather than generalizable findings that can advance the field. This paper illustrates why *general principles* of the translational process enhance mutual exchange between linguistic inquiry, model development, and application research (illustrated in Figure 1), and are key drivers of NLP advances.

We present a conceptual framework for Translational NLP, with specific elements of the translational process that are key to successful applications, each of which presents distinct areas for research. Our framework provides a concrete path for designing use-inspired basic research so that research products can effectively be turned into practical technologies, and provides the tools to understand why a technology translation succeeds or fails. A translational perspective further enables factorizing “grand challenge” research questions into clearly-defined pieces, producing intermediate results and driving new basic research questions. Our paper makes the following contributions:

- We characterize the stakeholders involved in the process of translating basic NLP advances to applications, and identify the roles they play in identifying new research problems (§3.1).
- We present a general-purpose checklist to use as a starting point for the translational process, to help integrate basic NLP innovations into applications and to identify basic research opportunities arising from application needs (§3.2).
- We present a case study in the medical domain illustrating how the elements of our Translational NLP framework can lead to new challenges for basic, applied, and translational NLP research (§4).

2 Defining Translational NLP

2.1 A third type of research

A long history of distinguishing between basic and applied research (Bush, 1945; Shneiderman, 2016) has noted that these terms are often relative; one researcher’s basic study is the application of another’s theory. In practice, basic and applied research in NLP are endpoints of a spectrum, rather than discrete categories. As *use-inspired* research, most NLP studies incorporate elements of both basic and applied research. We therefore define our key terms for this paper as follows:

Basic research Basic NLP research is focused on universal principles: linguistically-motivated study that guides model design (e.g., Recasens and Hovy (2009) for coreference, Kouloumpis et al. (2011) for sentiment analysis), or modeling techniques designed for general use across different settings and genres. Basic research tends to focus

on one problem at a time, and frequently leverages established datasets to provide a well-controlled environment for varying model design. Basic NLP research is intended to take the long view: it takes the time to investigate fundamental questions that may yield rewards for years to come.

Applied research Applied NLP research studies the intersection of universal principles with specific settings; it is responsive to the needs of commercial applications or researchers in other domains. Applied research utilizes real-world datasets, often specialized, and involves sources of noise and unreliability that complicate capturing linguistic regularities of interest. Applications often involve tackling multiple interrelated problems, and demand complex combinations of tools (e.g. using OCR followed by NLP to analyze scanned documents). Applied research is concrete and immediate, but may also be reactive and have a limited scope.

Translational research The term *translational* is used in medicine to describe research that aims to transform advances in basic knowledge (biological or clinical) to applications to human health (Butte, 2008; Rubio et al., 2010). Translational research is a distinct discipline bridging basic science and applications (Pober et al., 2001; Reis et al., 2010). We adopt the term *Translational NLP* to describe research bridging the gap between basic and applied NLP research, and aiming to understand the processes by which each informs the other. Section 4 presents one in-depth example; other salient examples include comparing the efficacy of domain adaptation methods for different application domains (Naik et al., 2019) and developing reusable software for processing specific text genres (Neumann et al., 2019). Translational research occupies a middle ground in the timeframe and complexity of solutions: it develops processes to rapidly and effectively integrate new innovations into applications to address emerging needs, and facilitates integration between pipelines of NLP tools.

2.2 Translation is bidirectional

In addition to “forward” motion of basic innovations into practical applications, the needs of real-world applications also provide significant opportunities for new fundamental research. Shneiderman’s model of “two parents, three children” (Shneiderman, 2016) provides an informative picture: combining a practical problem and a theoretical model yields (1) a solution to the problem,

(2) a refinement of the theory, and (3) guidance for future research. Tight links between basic research and applications have driven many major advances in NLP, from machine translation and dialog systems to search engines and question answering. Designing research with application needs in mind is a key impact criterion for both funding agencies (Christianson et al., 2018) and industry (Spector et al., 2012), and helps to identify new, high-impact research problems (Shneiderman, 2018).

2.3 NLP as a translational field: a historical perspective

The NLP field has always lain at the nexus of basic and applied research. Application needs have driven some of the most fundamental developments in the field, leading to explosions in basic research in new topics and on long-standing challenges.

The need to automatically translate Russian scientific papers in the early years of the Cold War led to some of the earliest NLP research, creating the still-thriving field of machine translation (Slocum, 1985). Machine translation has since helped drive many significant advances in basic NLP research, from the adoption of statistical models in the 1980s (Dorr et al., 1999) to neural sequence-to-sequence modeling (Sutskever et al., 2014) and attention mechanisms (Bahdanau et al., 2015).

Similarly, the rapid growth of the World Wide Web in the 1990s created an acute need for technologies to search the growing sea of information, leading to the development of NLP-based search engines such as Lycos (Mauldin, 1997), followed by PageRank (Page et al., 1999) and the growth of Google. The need to index and monetize vast quantities of textual information led to an explosion in information retrieval research, and the NLP field and ever-growing web data continue to co-develop.

In a more recent example, IBM identified automated question answering (QA) as a new business opportunity in a high-information world, and developed the Watson project (Ferrucci et al., 2010). Watson's early successes catapulted QA into the center of NLP research, where it has continued to drive both novel technology development and benchmark evaluation datasets used in hundreds of basic NLP studies (Rajpurkar et al., 2016).

These and other examples illustrate the key role that application needs have played in driving innovation in NLP research. This reflects not only the history of the field but the role that integrating basic

and applied research has in enriching scientific endeavor (Stokes, 1997; Branscomb, 1999; Narayana-murti et al., 2013; Shneiderman, 2016). An integrated approach has been cited by both Google (Spector et al., 2012) and IBM (McQueeney, 2003) as central to their successes in both business and research. The aim of our paper is to facilitate this integration in NLP more broadly, through presenting a rubric for studying and facilitating the process of getting both to and back from application.

2.4 A practical definition

For an operational definition of Translational NLP, it is instructive to consider four phases of a generic workflow for tackling a novel NLP problem using supervised machine learning.¹ First, a team of NLP experts works with subject matter experts (SMEs) to identify appropriate corpora, define concepts to be extracted, and construct annotation guidelines for the target task. Second, SMEs use these guidelines to annotate natural language data, using iterative evaluation, revision of guidelines, and re-annotation to converge on a high-quality *gold standard* set of annotations. Third, NLP experts use these annotations to train and evaluate candidate models of the task, joined with SMEs in a feedback loop to discuss results and needed revisions of goals, guidelines, and gold standards. Finally, buy-in is sought from SMEs and practitioners in the target domain, in a dialogue informed by empirical results and conceptual training. NLP adoption in practice identifies failure cases and new information needs, and the process begins again.

This laborious process is needed because of the gaps between expertise in NLP technology and expertise in use cases where NLP is applied. NLP expertise is needed to properly formulate problems, and subsequently to develop sound and generalizable solutions to those problems. However, for uptake (and therefore impact) to occur, these solutions must be based in deep expertise in the use case domain, reified in a computable manner through annotation or knowledge resource development. These distinct forms of expertise are generally found in different groups of individuals with complementary perspectives (see e.g. Kruschwitz and Hull (2017)).

Given this gap, we define **Translational NLP** as the development of theories, tools, and processes to enable the direct application of advanced NLP

¹While workflows will vary for different classes of NLP problems, dialogue between NLP experts and subject matter experts is at the heart of developing almost all NLP solutions.

tools in specific use cases. Implementing these tools and processes, and engaging with basic NLP experts and SMEs in their use, is the role of the Translational NLP scientist. Although every use case has unique characteristics, there are shared principles in designing NLP solutions that undergird the whole of the research and application process. These shared translational principles can be adopted by basic researchers to increase the impact of NLP methods innovations, and guide the translational researcher in developing novel efforts targeting fundamental gaps between basic research and applications. The framework presented in this paper identifies common variables and asks specific questions that can drive this research.

For examples of this process in practice, it is valuable to examine NLP development in the medical domain. Use-inspired NLP research has a long history in medicine (Sager et al., 1982; Ranum, 1989), frequently with an eye towards practical applications in research and care. Chapman et al. (2011) highlight shared tasks as a key step towards addressing numerous barriers to application of NLP on clinical notes, including lack of shared datasets, insufficient conventions and standards, limited reproducibility, and lack of user-centered design (all factors presenting basic research opportunities, in addition to NLP task improvement). Several efforts have explored the development of graphical user interfaces for conducting NLP tasks, including creation and execution of pipelines (Cunningham, 2002; D’Avolio et al., 2010, 2011; Soysal et al., 2018), although these efforts generally do not report on evaluation of usability by non-NLP experts. Usability has been investigated by other studies involving more focused tools aimed at specific NLP tasks, including concept searching (Hultman et al., 2018), annotation (Gobbel et al., 2014b,a), and interactive review of and update of text classification models (Trivedi et al., 2018, 2019; Savelka et al., 2015). Recent research has utilized interactive NLP tools for processing cancer research (Deng et al., 2019) and care (Yala et al., 2017) documents. By constructing, designing, and evaluating tools designed to simplify specific NLP processes, these efforts present examples of Translational NLP.

3 The Translational NLP framework

We present a conceptual framework for Translational NLP, to formalize shared principles describing how basic and applied research interact to cre-

ate NLP solutions. Our framework codifies fundamental variables in this process, providing a roadmap for negotiating the design of methodological innovations with an eye towards potential applications. Although it is certainly not the case that every basic research advance must be tied to a downstream application need, designing foundational technologies for potential application from the beginning produces more robust technologies that are easier to transfer to practical settings, increasing the impact of basic research. By defining common variables, our framework also provides a structure for aligning application needs to basic technologies, helping to identify potential failure points and new research needs early for faster adoption of basic NLP advances.

Our framework has two components:

1. A definition of broad classes of stakeholders in translating basic NLP innovations into applications, including the roles that each stakeholder plays in defining and guiding research;
2. A checklist of fundamental questions to structure the Translational NLP process, and to guide identification of basic research opportunities in specific application cases.

3.1 Stakeholders

NLP applications involve three broad categories of stakeholders, illustrated in Figure 2. Each contributes differently to technology implementation and identifying new research challenges.

NLP Experts NLP researchers bring key analytic skills to enable achieving the goals of an applied system. NLP experts provide *methodological sophistication* in models and paradigms for analyzing language, and an understanding of the nature of language and how it captures information. NLP researchers provide much-needed *data expertise*, including skills in obtaining, cleaning, and formatting data for machine learning and evaluation, as well as *conceptual models for representing information needs*. NLP scientists identify research opportunities in modeling information needs, bringing linguistic knowledge into the equation, and developing appropriate tools for application and reuse.

Subject Matter Experts Subject matter experts (SMEs) provide the context that helps to determine what information is important to analyze and what the outputs of applied NLP systems mean for the application setting. SMEs, from medical practitioners to legal scholars and financial experts, bring

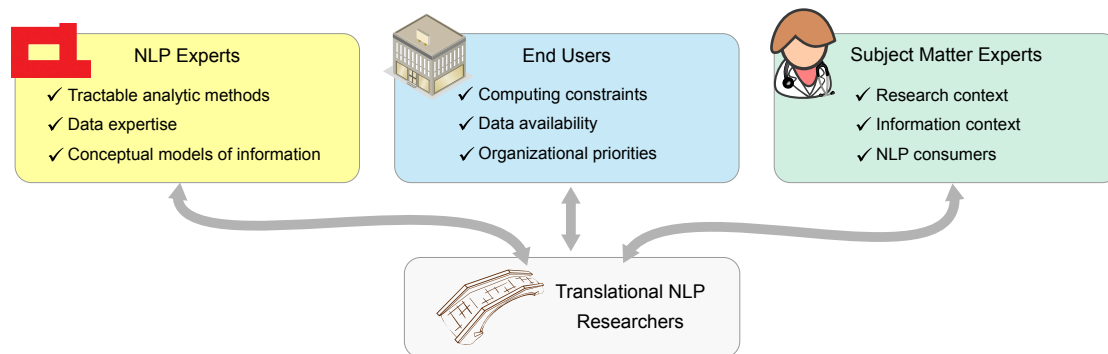


Figure 2: Attributes of key stakeholders in the translational process for NLP.

an understanding of *where relevant information can be found* (e.g., document sources (Fisher et al., 2016) and sections (Afzal et al., 2018)), which can help identify new types of language for basic researchers to study (Burstein, 2009; Crossley et al., 2014) and new challenges such as sparse complex information (Newman-Griffis and Fosler-Lussier, 2019) and higher-level structure in complex documents (Naik et al., 2019). In addition, the context that domain experts offer in terms of the needs of target applications feeds back into *evaluation* methods in the basic research setting (Graham, 2015).

SMEs are also the consumers of NLP solutions, as tools for their own research and applications. Thus, SMEs must also be consultants regarding the trustworthiness and reliability of proposed solutions, and can identify key application-specific concerns such as security requirements.

End Users The end users of NLP solutions span a range of roles, environmental contexts, and goals, each of which guides implementation factors of NLP applications. For example, collecting patient language in a lab setting, in a clinic, or at home will pose different challenges in each setting, which can inform the development of basic NLP methods. Application settings may have limited *computational resources*, motivating the development of efficient alternatives to high-resource models (e.g. Wang et al. (2020)), and have different *human factors* affecting information collection and use.

End users have different constraints on *data availability*, in terms of how much data of what types can be obtained from whom; the extensive work funded by DARPA’s Low Resource Languages for Emergent Incidents (LORELEI) initiative (Christianson et al., 2018) is a testament to the basic research arising from these constraints.

Beyond the individual domain expert, end users

use NLP technologies to address their own information needs according to the *priorities* of their organizations. These organizational priorities may conflict with existing modeling assumptions, highlighting new opportunities for basic research to expand model capabilities. For example, Shah et al. (2019) highlight the conceptual gap between predictive model performance in medicine and clinical utility to call for new research on utility-driven model evaluation. Spector et al. (2012) make a similar point about Google’s mission-driven research identifying unseen gaps for new basic research.

The role of the **Translational NLP** researcher is to interface with each of these stakeholders, to connect their goals, constraints, and contributions into a single applied system, and to identify new research opportunities where parts of this system conflict with one another. Notably, this creates an opportunity for valuable study of SME and end user research practices, and for participatory design of NLP research (Lazar et al., 2017). Our checklist, introduced in the next section, provides a structured framework for this translational process.

3.2 Translational NLP Checklist

The path between basic research and applications is often nebulous in NLP, limiting the downstream impact of modeling innovations and obscuring basic research challenges found in application settings. We present a general-purpose checklist covering fundamental variables in translating basic research into applications, which breaks down the translational process into discrete pieces for negotiation, measurement, and identification of new research opportunities. Our checklist, illustrated in Figure 3, is loosely ordered from initial design to application details. In practice, these items reflect different elements of the application process and are constantly




Figure 3: The eight items in our Translational NLP checklist, with key questions for each. Items are loosely ordered from initial design to application details, but should be regularly revisited in a feedback loop between application stakeholders.

re-evaluated via a feedback loop between the application stakeholders. While many of these items will be familiar to NLP researchers, each represents potential points of failure in translation. Designing the research process with these variables in mind will produce basic innovations that are more easily adopted for application and more directly connected to the challenges of real-world use cases.

We illustrate our items for two example cases:

Ex. 1: Analysis of multimodal clinical data (scanned text, tables, images) for patient diagnosis.

Ex. 2: Comparison of medical observations to government treatment and billing guidelines.

Information Need The initial step that guides an application is defining inputs and outputs, at two levels: (1) the *overall problem* to address with NLP (led by the subject matter expert), and (2) the *formal representation* of that problem (led by the NLP expert). The overall goal (e.g., “extract information on cancer from clinical notes”) determines the requirements of the solution, and is central to identifying a measurement of its effectiveness. Once the overall goal is determined, the next step is a formal representation of that goal in terms of text units (documents, spans) to analyze and what the analysis should produce (class labels, sequence annotations, document rankings, etc.). These requirements are tailored to specific applications and may not reflect standardized NLP tasks. For ex-

ample, a clinician interested in the documented reasoning behind a series of laboratory test orders needs: (1) the orders themselves (text spans); (2) the temporal sequence of the orders; and (3) a text span containing the justification for each order.

Ex. 1: type, severity, history of symptoms.

Ex. 2: clinical findings, logical criteria.

Data Characteristics A clear description of the *language data to be analyzed* is key to identifying appropriate NLP technologies. Data characteristics include the natural language(s) used (e.g., English, Chinese), the genre(s) of language to analyze (e.g., scientific abstracts, quarterly earnings reports, tweets, conversations), and the type(s) of linguistic community that produced them (e.g., medical practitioners, educators, policy experts). This information identifies the sublanguage(s) of interest (Grishman and Kittredge, 1986), which determine the availability and development of appropriate NLP tools (Grishman, 2001). Corporate disclosures, financial news reports, and tweets all require different processing strategies (Xing et al., 2018), as do tweets written by different communities (Blodgett et al., 2016; Groenwold et al., 2020).

Ex. 1: clinical texts, lab reports.

Ex. 2: clinical texts, legal guidelines.

Task Paradigms To address the overall goal with an NLP solution, it must be formulated in terms of one or more *well-defined NLP problems*. Many real-world application needs do not clearly correspond to a single benchmark task formulation. For example, our earlier example of the sequence of lab order justifications can be formulated as a sequence of: (1) Named Entity Recognition (treating the order types as named entities in a medical knowledge base); (2) time expression extraction and normalization; (3) event ordering; and (4) evidence identification. Breaking the application need into well-studied subproblems at design time enables faster identification and development of relevant NLP technologies, and highlights any portions of the goal that do *not* correspond with a known problem, requiring novel basic research.

Ex. 1: document type classification, OCR, information extraction (IE), patient classification.

Ex. 2: IE, natural language inference.

Available Resources The question of resources to support an NLP solution includes two distinct concerns: (1) *knowledge sources* available to represent salient aspects of the target task; and (2) *compute infrastructure* for NLP system execution

and deployment. Knowledge sources may be symbolic, such as knowledge graphs or gazetteers, or representational, such as representative corpora or pretrained language models. For some applications, powerful knowledge sources may be available (such as the UMLS (Bodenreider, 2004) for biomedical reasoning), while others are severely under-resourced (such as emerging geopolitical events, which may lack even relevant social media text). These resources in turn affect the kinds of technologies that are appropriate to use.

In terms of infrastructure, NLP technologies are deployed on a wide variety of systems, from commercial data centers to mobile devices. Each setting presents constraints of *limited resources* and *throughput requirements* (Nityasya et al., 2020). An application environment with a high maximum resource load but low median availability is amenable to batch processing architectures or approaches with high pretraining cost and low test-time cost. Pretrained word representations (Mikolov et al., 2013; Pennington et al., 2014) and language models (Peters et al., 2018; Devlin et al., 2019) are one example of fundamental technologies that address such a need. Throughput requirements, i.e., how much language input needs to be analyzed in a fixed amount of time, often require engineering optimization for specific environments (Afshar et al., 2019), but the need for faster runtime computation has led to many advances in machine learning for NLP, such as variational autoencoders (Kingma and Welling, 2014) and the Transformer architecture (Vaswani et al., 2017).

Ex. 1: UMLS, high GPU compute.

Ex. 2: UMLS, guideline criteria, low compute.

NLP Technologies The interaction between task paradigms, data characteristics, and available resources helps to determine what types of *implementations* are appropriate to the task. Implementations can be further broken down into *representation technologies*, for mathematically representing the language units to be analyzed; *modeling architectures*, for capturing regularities within that language; and *optimization strategies* (when using machine learning), for efficiently estimating model parameters from data. In low-resource settings, highly parameterized models such as BERT may not be appropriate, while large-scale GPU server farms enable highly complex model architectures. When the overall goal is factorized into multiple NLP tasks, optimization often involves

joint or multi-task learning (Caruana, 1997).

Ex. 1: large language models, dictionary matching, OCR, multi-task learning.

Ex. 2: dictionary matching, small neural models.

Evaluation Once a solution has been designed, it must be evaluated in terms of both *the specific NLP problem(s)* and *the overall goal* of the application. Standardized NLP task formulations typically define benchmark metrics which can be used for evaluating the NLP components: F-1 and AUC for information extraction, MRR and NDCG for information retrieval, etc. The design of these metrics is its own extensive area of research (Jones and Galliers, 1996; Hirschman and Thompson, 1997; Graham, 2015), and even established evaluation methods may be constantly revised (Grishman and Sundheim, 1995). Critically for the translational researcher, some metrics may be preferred over others (e.g., precision over recall), and standardized evaluation metrics may not reflect the goals and needs of applications (Friedman and Hripcsak, 1998). Improvements on standardized evaluation metrics (such as increased AUC) may even obscure degradations in application-relevant performance measures (such as decreased process efficiency). Translational researchers thus have the opportunity to work with NLP experts and SMEs to identify or develop metrics that capture both the effectiveness of the NLP system and its contribution to the application's overall goal.

Ex. 1: F-1, patient outcomes.

Ex. 2: F-1, billing rates.

Interpretation Interpretability and analysis of NLP and other machine learning systems has been the focus of extensive research in recent years (Gilpin et al., 2018; Belinkov and Glass, 2019), with debate over what constitutes an interpretation (Rudin, 2019; Wiegrefe and Pinter, 2019) and development of broad-coverage software packages for ease of use (Nori et al., 2019). For the translational researcher, the first step is to engage with SMEs to determine *what constitutes an acceptable interpretation* of an NLP system's output in the application domain (which may be subject to specific legal or ethical requirements around accountability in decision-making processes). This leads to an iterative process, working with SMEs and NLP experts to identify appropriately interpretable models, or to identify the need for new basic research on interpretability within the target domain.

Ex. 1: Evidence identification, model audits.

Ex. 2: Criteria visualization, model audits.

Application Engineering Last but not least, the translational process must also be concerned with the *implementation* of NLP solutions, both in terms of the specific technologies used and how they can fit in to broader information processing pipelines. The development of general-purpose NLP architectures such as the Stanford CoreNLP Toolkit (Manning et al., 2014), spaCy (Honnibal and Montani, 2017), and AllenNLP (Gardner et al., 2018), as well as more targeted architectures such as the clinical NLP framework presented by Wen et al. (2019), provide well-engineered frameworks for implementing new technologies in a way that is easy for others to both adopt and adapt for use in their own pipelines. Standardized data exchange frameworks such as UIMA (Ferrucci and Lally, 2004) and JSON make implementations more modular and easier to wire together. Leveraging tools and frameworks like these, together with good software design principles, makes NLP tools both easier to apply downstream and easier for other researchers to incorporate into their own work.

Ex. 1: Multiple interoperable technologies.

Ex. 2: Single decision support tool.

3.2.1 Translating methodology advances into existing applications

While the checklist items can guide initial design of a new NLP solution, they are equally applicable for incorporating new basic NLP innovations into existing solutions. Any new innovation can be reviewed in terms of our checklist items to identify new requirements or constraints (e.g., higher computational cost, more intuitive interpretability measures). The translational researcher can then work with NLP experts, SMEs, and the end users to determine how to incorporate the new innovation into the existing solution.

4 Case Study: NLP for Disability Review

We illustrate our Translational NLP framework using our recent line of research on developing NLP tools to assist US Social Security Administration (SSA) officials in reviewing applications for disability benefits (Desmet et al., 2020). The goal of this effort was to help identify relevant pieces of medical evidence for making decisions about disability benefits, analyzing vast quantities of medical records collected during the review process.

The stakeholders in this setting included: NLP researchers (interested in developing generalizable

methods); subject matter experts in disability and rehabilitation; and SSA end users (limited computing, large data but strictly controlled, overall priorities of efficiency and accuracy).

The Translational NLP checklist for this setting is shown in Table 1. This combination of factors has led to several translational studies, including:

- Newman-Griffis et al. (2018) developed a low-resource entity embedding method for domains with minimal knowledge sources (lack of Available Resources).
- Newman-Griffis and Zirikly (2018) analyzed the data size and representativeness tradeoff for information extraction in domains lacking large corpora (Available Resources).
- Newman-Griffis and Fosler-Lussier (2019) developed a flexible method for identifying sparse health information that is syntactically complex (challenging Data Characteristics).
- Newman-Griffis and Fosler-Lussier (2021) compared the Task Paradigms of classification and candidate selection paradigms for medical coding in a new domain.

While these studies do not systematically explore Evaluation, Interpretation, or Application Engineering, they illustrate how the characteristics of one application setting can lead to a line of Translational NLP research with broader implications. Several further challenges of this application area remain unstudied: for example, representing and modeling the complex timelines of persons with chronic health conditions and intermittent health care and adapting NLP systems to highly variable medical language from practitioners and patients around the US. These present intriguing challenges for basic NLP research that can inform many other applications beyond this case study.

Of course, these studies are far from the only examples of Translational NLP research. Many studies tackle translational questions, from domain adaptation (shifts in Data Characteristics) and low-resource learning (limited Available Resources), and the growing NLP literature in domain-specific venues such as medical research, law, finance, and more involves all aspects of the translational process. Rather, this case study is simply one illustration of how an explicitly translational perspective in study design can identify and connect broad opportunities for contributions to NLP research.

Information Need	<i>Overall goal:</i> Improve disability benefits review process by highlighting relevant information <i>Formal representation:</i> Spans of evidence, with attributes for activity type and level of limitation
Data Characteristics	Medical records and administrative forms from USA, mostly English
Task Paradigms	Information extraction (spans), information retrieval (documents), span classification (activity and limitations)
Available Resources	Minimal knowledge sources for function and disability, no public corpora; US government computing systems; high throughput requirements (thousands of records/day)
NLP Technologies	Low-latency, low-compute sequence models; rule-based systems
Evaluation	Standard metrics (F-1, accuracy). Information retrieval metrics reported for use case prototypes.
Interpretation	Interpretation needs primarily around human decision-making; NLP tools highlight and organize information in context. No ML interpretability reported in published results.
Application Engineering	Open-source implementations using standardized frameworks for preprocessing. No data exchange reported.

Table 1: Translational NLP checklist items for Disability Review case study, including notes on published results.

5 Discussion

Our paradigm of Translational NLP defines and gives structure to a valuable area of research not explicitly represented in the ACL community. We note that translational research is not meant to replace either basic or applied research, nor do we intend to say that all basic NLP studies must be tied to specific application needs. Rather we aim to highlight the value of studying the *processes* of turning basic innovations into successful applications. These processes, from scaling model computation to redesigning tools to meet changing application needs, can inform new research in model design, domain adaptation, etc., and can help us understand why some tools succeed in application while others fail. In addition to helping more innovations successfully translate, the principles outlined in this paper can be of use to basic and applied NLP researchers as well as translational ones, in identifying common variables and concerns to connect new work to the broader community.

Translational research is equally at home in industry and academia, and already occurring in both. While resource disparities between industrial and academic research increasingly push large-scale modeling efforts out of reach of academic teams, a translational lens can help to identify rich areas of knowledge-driven study that do not require exascale data or computing resources. The general principles and interdisciplinary nature of translational research make it a natural fit for public knowledge-driven academic settings, while its applicability to commercial needs is highly relevant to industry.

Our framework provides a starting point for the translational process, which will evolve differently

for every project. The specifics of different applications will expand our initial questions in different ways (e.g., “Data Characteristics” may involve multimodal data, or different language styles), and the dynamics of collaborations will shift answers over time (e.g., a change in evaluation criteria may motivate different model training approaches). Our checklist provides a minimal set of common questions, and can function as a touchstone for discussions throughout the research process, but it can and should be tailored to the nature of each project. Our framework is itself a preliminary characterization of Translational NLP research, and will evolve over time as the field continues to develop.

6 Conclusion

We have outlined a new model of NLP research, Translational NLP, which aims to bridge the gap between basic and applied NLP research with generalizable principles, tools, and processes. We identified key types of stakeholders in NLP applications and how they inform the translational process, and presented a checklist of common variables and translational principles to consider in basic, translational, or applied NLP research. The translational framework reflects the central role that integrating basic and applied research has played in the development of the NLP field, and is illustrated by both the broad successes of machine translation, speech processing, and web search, as well as many individual studies in the ACL community and beyond.

Acknowledgments

This work was supported by the National Library of Medicine of the National Institutes of Health

under award number T15 LM007059, and National Science Foundation grant 1822831.

References

- Majid Afshar, Dmitriy Dligach, Brihat Sharma, Xi-aoyuan Cai, Jason Boyda, Steven Birch, Daniel Valdez, Suzan Zelisko, Cara Joyce, François Modave, and Ron Price. 2019. [Development and application of a high throughput natural language processing architecture to convert all clinical documents in a clinical data warehouse into standardized medical vocabularies.](#) *Journal of the American Medical Informatics Association*, 26(11):1364–1369.
- Naveed Afzal, Vishnu Priya Mallipeddi, Sunghwan Sohn, Hongfang Liu, Rajeev Chaudhry, Christopher G Scott, Iftikhar J Kullo, and Adelaide M Arruda-Olson. 2018. [Natural language processing of clinical notes for identification of critical limb ischemia.](#) *International Journal of Medical Informatics*, 111:83–89.
- Laura J Anzaldi, Ashwini Davison, Cynthia M Boyd, Bruce Leff, and Hadi Kharrazi. 2017. [Comparing clinician descriptions of frailty and geriatric syndromes using electronic health records: a retrospective cohort study.](#) *BMC Geriatrics*, 17(1):248.
- Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. 2015. [Neural Machine Translation By Jointly Learning To Align and Translate.](#) In *ICLR 2015*.
- Yonatan Belinkov and James Glass. 2019. [Analysis methods in neural language processing: A survey.](#) *Transactions of the Association for Computational Linguistics*, 7:49–72.
- Su Lin Blodgett, Lisa Green, and Brendan O’Connor. 2016. [Demographic Dialectal Variation in Social Media: A Case Study of African-American English.](#) In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 1119–1130, Austin, Texas. Association for Computational Linguistics.
- Olivier Bodenreider. 2004. [The Unified Medical Language System \(UMLS\): integrating biomedical terminology.](#) *Nucleic Acids Research*, 32(90001):D267–D270.
- Lewis M Branscomb. 1999. [The false dichotomy: Scientific creativity and utility.](#) *Issues in Science and Technology*, 16(1):66–72.
- Jill Burstein. 2009. [Opportunities for Natural Language Processing Research in Education.](#) In *Computational Linguistics and Intelligent Text Processing. CICLing 2009*, pages 6–27, Berlin, Heidelberg. Springer Berlin Heidelberg.
- Vannevar Bush. 1945. *Science, The Endless Frontier*. U.S. Government Printing Office.
- Atul J Butte. 2008. [Translational Bioinformatics: Coming of Age.](#) *Journal of the American Medical Informatics Association*, 15(6):709–714.
- Rich Caruana. 1997. [Multitask learning.](#) *Machine learning*, 28(1):41–75.
- Wendy W Chapman, Prakash M Nadkarni, Lynette Hirschman, Leonard W D’Avolio, Guergana K Savova, and Ozlem Uzuner. 2011. [Overcoming barriers to NLP for clinical text: the role of shared tasks and the need for additional creative solutions.](#) *Journal of the American Medical Informatics Association*, 18(5):540–543.
- Caitlin Christianson, Jason Duncan, and Boyan Onyshkevych. 2018. [Overview of the DARPA LORELEI Program.](#) *Machine Translation*, 32(1):3–9.
- Scott A Crossley, Laura K Allen, Kristopher Kyle, and Danielle S McNamara. 2014. [Analyzing Discourse Processing Using a Simple Natural Language Processing Tool.](#) *Discourse Processes*, 51(5-6):511–534.
- Hamish Cunningham. 2002. [GATE, a General Architecture for Text Engineering.](#) *Computers and the Humanities*, 36(2):223–254.
- Leonard W. D’Avolio, Thien M. Nguyen, Wildon R. Farwell, Yongming Chen, Felicia Fitzmeyer, Owen M. Harris, and Louis D. Fiore. 2010. [Evaluation of a generalizable approach to clinical information retrieval using the automated retrieval console \(ARC\).](#) *Journal of the American Medical Informatics Association*, 17(4):375–382.
- Leonard W. D’Avolio, Thien M. Nguyen, Sergey Goryachev, and Louis D. Fiore. 2011. [Automated concept-level information extraction to reduce the need for custom software and rules development.](#) *Journal of the American Medical Informatics Association*, 18(5):607–613.
- Zhengyi Deng, Kanhua Yin, Yujia Bao, Victor Diego Armengol, Cathy Wang, Ankur Tiwari, Regina Barzilay, Giovanni Parmigiani, Danielle Braun, and Kevin S. Hughes. 2019. [Validation of a semiautomated natural language processing–based procedure for meta-analysis of cancer susceptibility gene penetrance.](#) *JCO Clinical Cancer Informatics*, (3):1–9.
- Shrey Desai, Geoffrey Goh, Arun Babu, and Ahmed Aly. 2020. [Lightweight convolutional representations for on-device natural language processing.](#) *arXiv preprint arXiv:2002.01535*.
- Bart Desmet, Julia Porcino, Ayah Zirikly, Denis Newman-Griffis, Guy Divita, and Elizabeth Rasch. 2020. [Development of natural language processing tools to support determination of federal disability benefits in the U.S.](#) In *Proceedings of the 1st Workshop on Language Technologies for Government and Public Administration (LT4Gov)*, pages 1–6, Marseille, France. European Language Resources Association.

- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [BERT: Pre-training of deep bidirectional transformers for language understanding](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.
- Bonnie J Dorr, Pamela W Jordan, and John W Benoit. 1999. [A Survey of Current Paradigms in Machine Translation](#). volume 49, pages 1–68. Elsevier.
- Hady Elsahar and Matthias Gallé. 2019. [To annotate or not? predicting performance drop under domain shift](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 2163–2173, Hong Kong, China. Association for Computational Linguistics.
- David Ferrucci, Eric Brown, Jennifer Chu-Carroll, James Fan, David Gondek, Aditya A Kalyanpur, Adam Lally, J William Murdock, Eric Nyberg, John Prager, and Others. 2010. Building Watson: An overview of the DeepQA project. *AI magazine*, 31(3):59–79.
- David Ferrucci and Adam Lally. 2004. UIMA: an architectural approach to unstructured information processing in the corporate research environment. *Natural Language Engineering*, 10:1–26.
- Ingrid E Fisher, Margaret R Garnsey, and Mark E Hughes. 2016. [Natural Language Processing in Accounting, Auditing and Finance: A Synthesis of the Literature with a Roadmap for Future Research](#). *Intelligent Systems in Accounting, Finance and Management*, 23(3):157–214.
- C Friedman and G Hripesak. 1998. Evaluating natural language processors in the clinical domain. *Methods of information in medicine*, 37(4-5):334.
- Matt Gardner, Joel Grus, Mark Neumann, Oyvind Tafjord, Pradeep Dasigi, Nelson F. Liu, Matthew Peters, Michael Schmitz, and Luke Zettlemoyer. 2018. [AllenNLP: A deep semantic natural language processing platform](#). In *Proceedings of Workshop for NLP Open Source Software (NLP-OSS)*, pages 1–6, Melbourne, Australia. Association for Computational Linguistics.
- L. H. Gilpin, D. Bau, B. Z. Yuan, A. Bajwa, M. Specter, and L. Kagal. 2018. [Explaining explanations: An overview of interpretability of machine learning](#). In *2018 IEEE 5th International Conference on Data Science and Advanced Analytics (DSAA)*, pages 80–89.
- Glenn T. Gobbel, Jennifer Garvin, Ruth Reeves, Robert M. Cronin, Julia Heavirland, Jenifer Williams, Allison Weaver, Shrimalini Jayaramaraja, Dario Giuse, Theodore Speroff, Steven H. Brown, Hua Xu, and Michael E. Matheny. 2014a. [Assisted annotation of medical free text using RapTAT](#). *Journal of the American Medical Informatics Association*, 21(5):833–841.
- Glenn T. Gobbel, Ruth Reeves, Shrimalini Jayaramaraja, Dario Giuse, Theodore Speroff, Steven H. Brown, Peter L. Elkin, and Michael E. Matheny. 2014b. [Development and evaluation of RapTAT: a machine learning system for concept mapping of phrases from medical narratives](#). *Journal of Biomedical Informatics*, 48:54–65.
- Yvette Graham. 2015. [Re-evaluating Automatic Summarization with BLEU and 192 Shades of ROUGE](#). In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pages 128–137, Lisbon, Portugal. Association for Computational Linguistics.
- Ralph Grishman. 2001. [Adaptive information extraction and sublanguage analysis](#). In *Proceedings of the Workshop on Adaptive Text Extraction and Mining, Seventeenth International Joint Conference on Artificial Intelligence (IJCAI-2001)*, pages 1–4, Seattle, Washington, USA.
- Ralph Grishman and R Kittredge. 1986. *Analyzing language in restricted domains: Sublanguage description and processing*. Lawrence Erlbaum Associates.
- Ralph Grishman and Beth Sundheim. 1995. Design of the MUC-6 evaluation. In *Proceedings of the 6th Conference on Message Understanding*, pages 1–11. Association for Computational Linguistics.
- Sophie Groenwold, Lily Ou, Aesha Parekh, Samhita Honnavalli, Sharon Levy, Diba Mirza, and William Yang Wang. 2020. [Investigating African-American Vernacular English in Transformer-Based Text Generation](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 5877–5883, Online. Association for Computational Linguistics.
- Lynette Hirschman and Henry S. Thompson. 1997. Overview of evaluation in speech and natural language processing. In Ronald A. Cole, editor, *Survey of the State of the Art in Human Language Technology*. Cambridge University Press.
- Matthew Honnibal and Ines Montani. 2017. [spacy 2: Natural language understanding with bloom embeddings, convolutional neural networks and incremental parsing](#).
- Gretchen Hultman, Reed McEwan, Serguei Pakhomov, Elizabeth Lindemann, Steven Skube, and Genevieve B. Melton. 2018. Usability Evaluation of an Unstructured Clinical Document Query Tool for Researchers. *AMIA Joint Summits on Translational Science proceedings. AMIA Joint Summits on Translational Science*, 2017:84–93.

- Karen Spärck Jones and Julia R. Galliers. 1996. *Evaluating Natural Language Processing Systems: An Analysis and Review*. Springer-Verlag, Berlin, Heidelberg.
- Diederik P Kingma and Max Welling. 2014. Auto-encoding variational bayes. *ICLR 2014*.
- Efthymios Kouloumpis, Theresa Wilson, and Johanna Moore. 2011. Twitter sentiment analysis: The good the bad and the omg! In *Fifth International AAAI conference on weblogs and social media*. Citeseer.
- Udo Kruschwitz and Charlie Hull. 2017. [Searching the Enterprise](#). *Foundations and Trends in Information Retrieval*, 11(1):18.
- Jonathan Lazar, Jinjuan Heidi Feng, and Harry Hochheiser. 2017. *Research methods in human-computer interaction*. Morgan Kaufmann.
- Christopher Manning, Mihai Surdeanu, John Bauer, Jenny Finkel, Steven Bethard, and David McClosky. 2014. [The Stanford CoreNLP natural language processing toolkit](#). In *Proceedings of 52nd Annual Meeting of the Association for Computational Linguistics: System Demonstrations*, pages 55–60, Baltimore, Maryland. Association for Computational Linguistics.
- M I Mauldin. 1997. [Lycos: design choices in an Internet search service](#). *IEEE Expert*, 12(1):8–11.
- David F McQueeney. 2003. IBM’s evolving research strategy. *Research-Technology Management*, 46(4):20–27.
- Tomas Mikolov, Kai Chen, Greg Corrado, and Jeffrey Dean. 2013. [Efficient Estimation of Word Representations in Vector Space](#). *arXiv preprint arXiv:1301.3781*, pages 1–12.
- Aakanksha Naik, Luke Breittfeller, and Carolyn Rose. 2019. [TDDiscourse: A Dataset for Discourse-Level Temporal Ordering of Events](#). In *Proceedings of the 20th Annual SIGdial Meeting on Discourse and Dialogue*, pages 239–249, Stockholm, Sweden. Association for Computational Linguistics.
- Venkatesh Narayanamurti, Tolu Odumosu, and Lee Vinsel. 2013. RIP: The basic/applied research dichotomy. *Issues in Science and Technology*, 29(2):31–36.
- Mark Neumann, Daniel King, Iz Beltagy, and Waleed Ammar. 2019. [ScispaCy: Fast and robust models for biomedical natural language processing](#). In *Proceedings of the 18th BioNLP Workshop and Shared Task*, pages 319–327, Florence, Italy. Association for Computational Linguistics.
- Denis Newman-Griffis and Eric Fosler-Lussier. 2019. [HARE: a Flexible Highlighting Annotator for Ranking and Exploration](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP): System Demonstrations*, pages 85–90, Hong Kong, China. Association for Computational Linguistics.
- Denis Newman-Griffis and Eric Fosler-Lussier. 2021. [Automated Coding of Under-Studied Medical Concept Domains: Linking Physical Activity Reports to the International Classification of Functioning, Disability, and Health](#). *Frontiers in Digital Health*, 3:620828.
- Denis Newman-Griffis, Albert M Lai, and Eric Fosler-Lussier. 2018. [Jointly embedding entities and text with distant supervision](#). In *Proceedings of The Third Workshop on Representation Learning for NLP*, pages 195–206, Melbourne, Australia. Association for Computational Linguistics.
- Denis Newman-Griffis and Ayah Ziriky. 2018. [Embedding transfer for low-resource medical named entity recognition: A case study on patient mobility](#). In *Proceedings of the BioNLP 2018 workshop*, pages 1–11, Melbourne, Australia. Association for Computational Linguistics.
- Made Nindyatama Nityasya, Haryo Akbarianto Wibowo, Radityo Eko Prasajo, and Alham Fikri Aji. 2020. No Budget? Don’t Flex! Cost Consideration when Planning to Adopt NLP for Your Business. *arXiv preprint arXiv:2012.08958*.
- Harsha Nori, Samuel Jenkins, Paul Koch, and Rich Caruana. 2019. [InterpretML: A unified framework for machine learning interpretability](#). *arXiv preprint arXiv:1909.09223*.
- Lawrence Page, Sergey Brin, Rajeev Motwani, and Terry Winograd. 1999. The PageRank citation ranking: Bringing order to the web. Technical report, Stanford InfoLab.
- Jeffrey Pennington, Richard Socher, and Christopher Manning. 2014. [GloVe: Global vectors for word representation](#). In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1532–1543, Doha, Qatar. Association for Computational Linguistics.
- Matthew Peters, Mark Neumann, Mohit Iyyer, Matt Gardner, Christopher Clark, Kenton Lee, and Luke Zettlemoyer. 2018. [Deep contextualized word representations](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 2227–2237, New Orleans, Louisiana. Association for Computational Linguistics.
- Jordan S Pober, Crystal S Neuhauser, and Jeremy M Pober. 2001. Obstacles facing translational research in academic medical centers. *The FASEB Journal*, 15(13):2303–2313.
- Protiva Rahman, Arnab Nandi, and Courtney Hebert. 2020. [Amplifying domain expertise in clinical data pipelines](#). *JMIR Med Inform*, 8(11):e19612.

- Pranav Rajpurkar, Jian Zhang, Konstantin Lopyrev, and Percy Liang. 2016. SQuAD: 100,000+ Questions for Machine Comprehension of Text. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 2383–2392.
- David L Ranum. 1989. Knowledge-based understanding of radiology text. *Computer methods and programs in biomedicine*, 30(2-3):209–215.
- Marta Recasens and Eduard Hovy. 2009. [A deeper look into features for coreference resolution](#). In *Proceedings of the 7th Discourse Anaphora and Anaphor Resolution Colloquium on Anaphora Processing and Applications*, DAARC '09, page 29–42, Berlin, Heidelberg. Springer-Verlag.
- Steven E Reis, Lars Berglund, Gordon R Bernard, Robert M Califf, Garret A Fitzgerald, Peter C Johnson, National Clinical Consortium, and Translational Science Awards. 2010. [Reengineering the national clinical and translational research enterprise: the strategic plan of the National Clinical and Translational Science Awards Consortium](#). *Academic Medicine*, 85(3):463–469.
- Doris McGartland Rubio, Ellie E Schoenbaum, Linda S Lee, David E Scheingart, Paul R Marantz, Karl E Anderson, Lauren Dewey Platt, Adriana Baez, and Karin Esposito. 2010. [Defining translational research: implications for training](#). *Academic Medicine*, 85(3):470–475.
- Cynthia Rudin. 2019. Stop explaining black box machine learning models for high stakes decisions and use interpretable models instead. *Nature Machine Intelligence*, 1(5):206–215.
- N Sager, IDJ Bross, G Story, P Bastedo, E Marsh, and D Shedd. 1982. Automatic encoding of clinical narrative. *Computers in Biology and Medicine*, 12(1):43–56.
- Jaromir Savelka, Gaurav Trivedi, and Kevin Ashley. 2015. Applying an interactive machine learning approach to statutory analysis. In *JURIX 2015 - the 28th International Conference on Legal Knowledge and Information Systems*.
- Nigam H Shah, Arnold Milstein, and PhD Bagley Steven C. 2019. [Making Machine Learning Models Clinically Useful](#). *JAMA*, 322(14):1351–1352.
- Ben Shneiderman. 2016. *The new ABCs of research: Achieving breakthrough collaborations*. Oxford University Press.
- Ben Shneiderman. 2018. [Twin-Win Model: A human-centered approach to research success](#). *Proceedings of the National Academy of Sciences of the United States of America*, 115(50):12590–12594.
- Stefan Slater, Srećko Joksimović, Vitomir Kovanovic, Ryan S Baker, and Dragan Gasevic. 2017. Tools for educational data mining: A review. *Journal of Educational and Behavioral Statistics*, 42(1):85–106.
- Jonathan Slocum. 1985. A Survey of Machine Translation: Its History, Current Status, and Future Prospects. *Comput. Linguist.*, 11(1):1–17.
- Ergin Soysal, Jingqi Wang, Min Jiang, Yonghui Wu, Serguei Pakhomov, Hongfang Liu, and Hua Xu. 2018. [CLAMP – a toolkit for efficiently building customized clinical natural language processing pipelines](#). *Journal of the American Medical Informatics Association*, 25(3):331–336. Publisher: Oxford Academic.
- Alfred Spector, Peter Norvig, and Slav Petrov. 2012. Google’s hybrid approach to research. *Communications of the ACM*, 55(7):34–37.
- Donald E Stokes. 1997. *Pasteur’s quadrant: Basic science and technological innovation*. Brookings Institution Press.
- Ilya Sutskever, Oriol Vinyals, and Quoc V Le. 2014. Sequence to sequence learning with neural networks. In *Advances in neural information processing systems*, pages 3104–3112.
- Gaurav Trivedi, Esmaeel R Dadashzadeh, Robert M Handzel, Wendy W Chapman, Shyam Visweswaran, and Harry Hochheiser. 2019. [Interactive NLP in Clinical Care: Identifying Incidental Findings in Radiology Reports](#). *Appl Clin Inform*, 10(04):655–669.
- Gaurav Trivedi, Phuong Pham, Wendy W. Chapman, Rebecca Hwa, Janyce Wiebe, and Harry Hochheiser. 2018. [NLPReViz: an interactive tool for natural language processing on clinical text](#). *Journal of the American Medical Informatics Association*, 25(1):81–87.
- 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*, pages 5998–6008.
- Hanrui Wang, Zhanghao Wu, Zhijian Liu, Han Cai, Ligeng Zhu, Chuang Gan, and Song Han. 2020. Hat: Hardware-aware transformers for efficient natural language processing. *arXiv preprint arXiv:2005.14187*.
- Andrew Wen, Sunyang Fu, Sungrim Moon, Mohamed El Wazir, Andrew Rosenbaum, Vinod C Kaggal, Sijia Liu, Sunghwan Sohn, Hongfang Liu, and Jungwei Fan. 2019. Desiderata for delivering nlp to accelerate healthcare ai advancement and a mayo clinic nlp-as-a-service implementation. *NPJ Digital Medicine*, 2(1):1–7.
- Sarah Wiegrefe and Yuval Pinter. 2019. [Attention is not not explanation](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 11–20, Hong Kong, China. Association for Computational Linguistics.

Frank Z Xing, Erik Cambria, and Roy E Welsch. 2018. [Natural language based financial forecasting: a survey](#). *Artificial Intelligence Review*, 50(1):49–73.

Adam Yala, Regina Barzilay, Laura Salama, Molly Griffin, Grace Sollender, Aditya Bardia, Constance Lehman, Julliette M Buckley, Suzanne B Coopey, Fernanda Polubriaginof, et al. 2017. Using machine learning to parse breast pathology reports. *Breast Cancer Research and Treatment*, 161(2):203–211.