

A Diachronic Analysis of Paradigm Shifts in NLP Research: When, How, and Why?

Aniket Pramanick¹, Yufang Hou², Saif M. Mohammad³, Iryna Gurevych¹

¹Ubiquitous Knowledge Processing Lab (UKP Lab)

Department of Computer Science and Hessian Center for AI (hessian.AI)

²IBM Research Europe, Ireland

³National Research Council Canada

www.ukp.tu-darmstadt.de, yhou@ie.ibm.com, saif.mohammad@nrc-cnrc.gc.ca

Abstract

Understanding the fundamental concepts and trends in a scientific field is crucial for keeping abreast of its continuous advancement. In this study, we propose a systematic framework for analyzing the evolution of research topics in a scientific field using causal discovery and inference techniques. We define three variables to encompass diverse facets of the evolution of research topics within NLP and utilize a causal discovery algorithm to unveil the causal connections among these variables using observational data. Subsequently, we leverage this structure to measure the intensity of these relationships. By conducting extensive experiments on the *ACL Anthology* corpus, we demonstrate that our framework effectively uncovers evolutionary trends and the underlying causes for a wide range of NLP research topics. Specifically, we show that tasks and methods are primary drivers of research in NLP, with datasets following, while metrics have minimal impact.¹

1 Introduction

Experts in a field sometimes conduct historical studies to synthesize and document the key research ideas, topics of interest, methods, and datasets that shaped a field of study. They document how new research topics eclipsed older ones and contributed to shaping the trajectory of the research area (Kuhn, 1970). Aspiring scientists learn the craft of their discipline by delving into the examination of past scientific accomplishments documented in research papers. However, conducting such a historical study is challenging: Experts in a field rely on years of experience and peruse large amounts of past published articles to determine the chronological progression of a research field. Further, the exponential growth of scientific publications in recent years has rendered it arduous even for domain experts to stay current. Therefore, an automated

method to track the temporal evolution of research topics can be beneficial in offering an overview of the field and assisting researchers in staying abreast of advancements more efficiently.

In this work, we propose a systematic framework to examine the evolutionary journey of research topics within the realm of Natural Language Processing (NLP), harnessing causal discovery and inference techniques. Prior research on historical analysis of NLP has predominantly concentrated on scrutinizing metadata associated with research papers (Hall et al., 2008; Mohammad, 2019; Uban et al., 2021; Singh et al., 2023; Wahle et al., 2023) such as number of citations, title, author profile, affiliation, and publication venue. These studies have examined the research trends through unigram or bigram frequency analysis, but they do not provide insights into the underlying causes propelling these research topics.

Our study centers on four distinct fundamental types of entities in NLP research: **tasks** representing well defined problems; **methods**, signifying the solutions or approaches employed to tackle the tasks; **datasets**, indicating the relevant textual resources such as corpora and lexicons; and **metrics**, encompassing the evaluation techniques tailored to specific tasks. We abbreviate these types as **TDMM** for short. Specifically, we examine the interplay between an NLP task that is commonly viewed as a focused research topic (e.g., *Machine Translation*) and the key *entities* that exert pivotal influence on the target task (such as “*BLEU*” (Papineni et al., 2002) or “*Transformers*” (Vaswani et al., 2017)).

Our goal is to identify the TDMM entities (E) associated with a specific task (t) and assess their causal influence on the task’s research trends (**TDMM-Task causal analysis**). Specifically, we address the following key research questions associated with a task entity t : (a) Which entities E effectively indicate the research trends for this task

¹We publish the code and dataset for our experiments at <https://github.com/UKPLab/CausalNLPtrends>




Figure 1: Evolution of Machine Translation (MT) research. Blue line: Number of MT papers (1979-2022). Tables show the top causal entities/types for different periods (excluding 1979-1989 due to limited MT papers).

t ? (b) Are there discernible causal relationships between t and E ? (c) What is the extent of the causal impact exerted by E on t ?

Unlike Uban et al. (2021) and Koch et al. (2021) that heavily rely on manual annotations and have limited coverage, our analysis is based on TDMM entities automatically extracted from 55K papers in the ACL Anthology². Our framework not only recognizes the key entities driving the research direction of a research topic but also measures the causal effects of these entities on the target topic in an end-to-end fashion. Figure 1 shows the most influential entities for *Machine Translation* (MT) in different time periods. For instance, “*statistical models*” used to be the popular method for MT in 1990-2002, and the evaluation metric “*BLEU*” is one of the top causal entities driving the MT research in 2003-2017. In the era of pre-trained large language models (LLMs) starting from 2018, “*transformer*” has become the popular method for MT. For another research topic of “*Speech recognition*”, our framework uncovers the influential role of “*language modeling*” between 1979 to 2022, where speech recognition models utilize probability scores from language models to recognize coherent text from speech (Negri et al., 2014).

In this work, we analyze 16 tasks from a diverse set of research areas identified by ACL 2018 organizers. Our framework is versatile and applicable to other tasks and domains, benefiting both young and experienced researchers. It can aid in litera-

ture surveys by identifying related research areas and enable young researchers to delve into new research focuses by establishing connections among different research areas.

In summary, we make *three-fold* contributions in this study: **Firstly**, we propose a framework to quantify research activities, including (1) trends and stability of an NLP research task, and (2) relation intensity between TDMM entities and NLP research tasks. **Secondly**, we employ causal analysis algorithms to uncover causal structures and measure effects between tasks and related TDMM entities (*TDMM-Task causal analysis*). To the best of our knowledge, this represents the first historical study of a scientific research anthology from a causal perspective. **Finally**, through extensive experiments on the ACL Anthology, we offer an empirical overview of the NLP research landscape. In the following sections, we will refer to *TDMM-Task causal analysis* as *causal analysis*.

2 Related Work

Scientific Trends Analysis The analysis of scientific trends has been a research focus since Hall et al. (2008). In the field of “scientometrics”, extensive literature explores citation patterns and utilizes topological measures in citation networks for trend analysis (Small, 2006; Shibata et al., 2008; Boyack and Klavans, 2022).

Another line of research focuses on metadata and content analysis. For instance, Prabhakaran et al. (2016) employed rhetorical framing to exam-

²<https://aclanthology.org/>

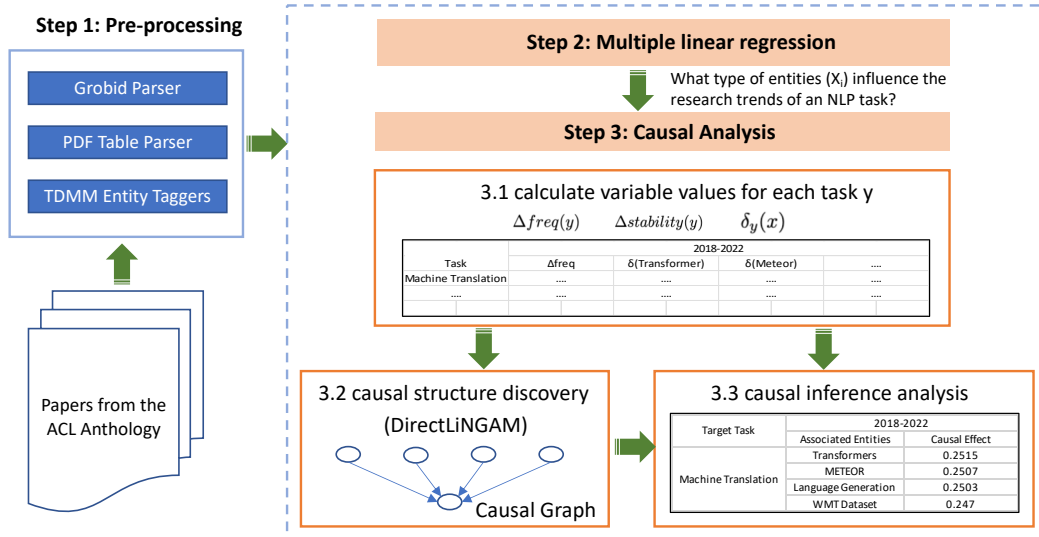


Figure 2: System architecture.

ine trend patterns. Grudin (2009), Liu et al. (2015), and Mohammad (2019) investigated the interaction between the topics in publications, research grants, author profiles, highly impactful papers, and dataset usage patterns. Additionally, Koch et al. (2021) studied dataset usage patterns among different research communities, while Uban et al. (2021) analyzed relationships between NLP research topics based on their co-occurrence in text and the degree of correlation between their popularity over time. In our work, we develop entity recognition models to extract TDMM entities from NLP research papers and focus on analyzing the causal relations between a task entity and its related TDMM entities.

Causality in NLP Existing works on NLP applying causal analysis algorithms mainly focus on two directions. The first line of work discovers causal relations among textual features or expressions of events in texts and uses them in various downstream tasks, such as question answering (Oh et al., 2016), commonsense reasoning (Bosselut et al., 2019; Sap et al., 2019), and relation extraction (Do et al., 2011; Mirza and Tonelli, 2014; Duni et al., 2017).

In another avenue of this field, researchers represent causal elements using textual features (Jin et al., 2021; Fong and Grimmer, 2016; Veitch et al., 2020; Keith et al., 2020) and define the causal graph structure based on domain knowledge. Our work falls within this line of research, where we employ causal algorithms to analyze the trends in NLP research topics and the underlying causes.

3 Data Collection

ACL Anthology Corpus Following prior work by Mohammad (2020), we utilize ACL Anthology as the source of NLP Research papers. For this work, we collect 55,366 NLP papers that belong to the “ACL Events” category³ from the ACL anthology published between 1979 and 2022. For each paper, we use GROBID (GRO, 2008–2022) and the PDF table parser from Hou et al. (2019) to extract sentences from each of the individual sections as well as from the table and figure captions. In a post-processing step, we remove all the URLs from the extracted sentences. On average, we have 1,258 papers per year and 1,117 sentences per paper.

It is worth noting that certain NLP paper preprints may become accessible on preprint servers before they are officially published in the ACL Anthology. However, we argue that the peer review process in ACL Anthology serves as a robust quality assurance mechanism. Hence, we consider ACL Anthology a more reliable source compared to preprint servers.

TDMM Entity Extraction To identify *tasks*, *datasets*, *metrics*, and *methods* entities from NLP papers, we developed two entity taggers based on Flair (Akbik et al., 2018). The first tagger is based on the TDMSci annotations (Hou et al., 2021) for recognizing *task*, *dataset*, and *metric* entities. The second tagger is trained using the SciERC dataset

³This category covers major NLP conferences, workshops, and journals including ACL, NAACL, EMNLP, EACL, AACL, CL, and TACL. Additionally, we also include papers published at COLING from the “Non-ACL events category”.

Period	Years	Key Research Themes
Early Years	1979–1989	Foundational work in syntactic parsing, machine translation, and information retrieval.
Formative Years	1990–2002	Advances in language modeling, named entity recognition, and discourse analysis (research focus shifted towards data-driven approaches).
Statistical Revolution & Neural Networks	2003–2017	Focus on statistical techniques (text classification, statistical machine translation, etc.) and resurgence of neural networks (word embeddings, neural machine translation, etc.)
Deep Learning Era	2018–2022	Dominance of transformer-based architectures (BERT and their variants).

Table 1: Chronological Periods of NLP Research.

(Luan et al., 2018) to extract *method* entities. On the testing datasets of TDMSci and SciERC, the two taggers achieve a micro-average F1 of 0.77 and 0.78 for the type partial match (Segura-Bedmar et al., 2013), respectively. In type partial match, a predicted entity is considered correct if it partially overlaps with a gold entity and has the same type. For example, “*Penn Treebank*” is counted as a correct prediction even if the corresponding gold annotation is “*Penn Treebank dataset*”.

To further improve the precision of the TDMM taggers, we include only entities that appear in more than five papers in the dataset. For each paper, we collect the most frequent task mentions appearing in the *title*, *abstract*, *experiment section*, *table*, and *figure captions* to approximate the tasks that the paper has done research on.

Taxonomy for Periods of Reference In order to facilitate in-depth analysis, in this paper, we adopt a taxonomy that partitions our reference time frame (1979-2022) into four distinct intervals. Table 1 illustrates the defined intervals. These intervals have been designed to approximate the overarching trends observed in NLP research throughout the years, aligning with our perspective on the field’s evolution. It is important to acknowledge that the exact boundaries and thematic emphases may differ based on varying perspectives and specific research areas within NLP. However, we highlight that our framework and methodologies are highly adaptable, allowing end users to effortlessly apply them to any desired time interval or a specific analysis.

4 Entity Influence in NLP Research: A Regression Analysis

Before conducting the causal analysis, we aim to identify the key variables that significantly impact the evolution of NLP Research. Specifically, we investigate which types of entities exert the most

Variables	R-Squared (\uparrow)
unique tasks	0.87
+ unique datasets	0.91
+ unique methods	0.93
+ unique metrics	0.97

Table 2: Variable Selection for Regression.

influence on the research direction of NLP. To achieve this understanding, we employ Multiple Linear Regression (see Appendix D for details), a widely utilized tool in economics research (Barrios and Hochberg, 2020). Figure 2 (step1/step2) illustrates the framework.

Our analysis assumes that if the TDMM entities have played a role in the emergence or disappearance of task entities, this influence will be reflected in the number of unique task entities in subsequent years, which can be captured through regression analysis. While the study does not provide specific information on the precise influence of each TDMM entity on individual task entities, the partial regression coefficients shed light on the types of entities responsible for influencing the overall task entity landscape.

Method. Mathematically, we predict the number of task entities Y^t in a given year t as a function of the cumulative counts of all types of entities $\{X_i^{t-1}\}$ (TDMM entities) until that year, $t-1$, given by $Y^t = r_0 + \sum_i r_i X_i^{t-1}$. $\{r_i\}$ quantifies the relationship strength between the predicted variable (number of task entities) and the independent variables (number of TDMM entities).

Evaluation. We evaluate the regression model using the R^2 measure (coefficient of determination) to assess the goodness of fit. Additionally, we perform a null hypothesis test to determine the statistical significance of the partial regression co-

efficients.

Results and Discussion.

1) Optimized Number of Variables. In our initial experiment, we determine the optimal number of variables and summarize the corresponding R^2 values in Table 2. Additionally, all regression coefficients are statistically significant at 5% level, indicating their strong relationship with the predicted variable. **Discussion:** The overall results indicate that the model achieves a good fit to the data when all four variables (number of tasks, datasets, metrics, and method entities) are used to predict the number of task entities in subsequent years. We also explore the possibility of reducing the number of variables while maintaining similar performance. Interestingly, using only one variable results in a significant drop of 0.1 in the R^2 value (R^2 value 0.87), indicating a poor fit to the model. Conversely, increasing the number of variables improves the model fit, suggesting the significance of all four variables in analyzing research trends (R^2 value 0.97). It is worth noting that we exhaustively explored various combinations of variables, including those presented in the table, and consistently obtained similar results.

2) Influence of the Variables. In the second experiment, we assess the association between the target variable and each independent variable. In Table 3, we present the regression coefficients corresponding to each entity type. Larger values of regression coefficients indicate a stronger relationship between the target variable and the respective independent variable. **Discussion:** Overall, we note that the gradual emergence of newer tasks has been a driving force behind research progress. However, when we analyze the trends within each year interval, we uncover more nuanced patterns. During the *Early Years (1979–1989)*, when NLP was in its nascent stage as an independent research field, the focus was on creating new *datasets* to fuel research advancements. In the *Formative Years (1990–2002)*, we witnessed the introduction of new *methods*, particularly data-driven approaches, which played a crucial role in shaping the field. Subsequently, *from 2003 to 2017*, *statistical methods* underwent a revolution, and later in the same period, *neural network methods* experienced a resurgence, indicating significant shifts in research trends. Now, in the present *Deep Learning Era (2018–2022)*, we observe a rapid creation of

Years	Partial Regression Coefficient			
	Tasks	Datasets	Methods	Metrics
1979–1989	0.35	2.24	0.21	0.02
1990–2002	0.82	0.89	2.86	0.81
2003–2017	5.37	6.26	7.00	0.69
2018–2022	1.47	3.38	1.79	0.41
1979 - 2022	3.50	1.07	2.92	0.54

Table 3: Variables Influencing NLP task entities.

newer *datasets* in a relatively short span of time, driven by the research needs and the data requirements of deep learning models. These highlight key factors influencing research trajectory over time.

5 Causal Methodology for NLP Research Analysis

Drawing on the insights gained from the Regression Analysis (Section 4), we now establish the cornerstone of our study by defining three causal variables that drive the causal analysis in the subsequent sections. Using causal discovery and inference techniques, we analyze the causal relationships among the variables and measure the impact of TDMM entities on target task entities based on these relationships. Figure 2 illustrates the architecture that underpins our framework.

5.1 Causal Variables

Task Frequency Shift Value: Distinguishing from the previous approaches (Tan et al., 2017; Prabhakaran et al., 2016), that rely on word frequencies, we define task frequency $f(y)_t$ as the number of published papers focusing on a specific task y in a given year t , normalized by the total number of papers published on the same year. The task frequency shift value $\Delta freq_{t_1}^{t_2}(y)$ captures the average change in the number of published papers on y between two years $t_1 < t_2$. This value serves as a measure of the research trends associated with the task during that time interval, indicating whether it experienced growth or decline. The frequency shift value is given by: $\Delta freq_{t_1}^{t_2}(y) = \frac{f(y)_{t_2} - f(y)_{t_1}}{t_2 - t_1}$.

Task Stability Value: We introduce the concept of task stability value to measure the change in the research context of a given task, y , between two years, $t_1 < t_2$. This value quantifies the overlap in neighboring TDMM entities that appear in the same publication as y within the specified time interval. To calculate task stability, we adapt

the semantic stability approach of [Wendlandt et al. \(2018\)](#) to our setting and define it specifically for task entities. Initially, we represent each paper in our dataset as a sequence of TDMM entity mentions, removing non-entity tokens. We then employ “Skip-gram with negative sampling” ([Mikolov et al., 2013](#)) to obtain embeddings from this representation. Formally, let e_1, e_2, \dots, e_n be this entity representation of a paper, and the objective of skip-gram is to maximize the mean log probability $\frac{1}{n} \sum_{i=1}^n \sum_{-c \leq j \leq c} \log p(e_{i+j}|e_i)$, where c is called the context window size. Finally, the task stability value $\Delta stability_{t_1}^{t_2}(y)$ of y between t_1 and t_2 is computed as the percentage overlap between the nearest l neighboring entities of the given task in two representation spaces. The stability value is given by: $\Delta stability_{t_1}^{t_2}(y) = \frac{|\mathcal{N}_{t_1}^l(y) \cap \mathcal{N}_{t_2}^l(y)|}{|\mathcal{N}_{t_1}^l(y) \cup \mathcal{N}_{t_2}^l(y)|}$, where $\mathcal{N}_t^l(y)$ is the set of l neighbours of y in the representation space of year t . In this study, we consider the context window c to encompass the entire document, and we set the value of l to 5.

Entity Change Value: We use entity change value to track emerging and disappearing of specific TDMM entities associated with a task, quantifying these changes and capturing related entity occurrences within a specific time period. Put simply, we measure the difference in the co-occurrence frequency of a TDMM entity x and a task y between two years t_1 and t_2 . When we identify a significant change in the co-occurrence frequency of x and y over this period, it likely signals a shift in the relation between x and y and, in turn, a shift in NLP Research trends. We define entity change value $\delta_y(x)_{t_1}^{t_2}$ of an entity x of type $\tau(x) \in \{\text{task, dataset, metric, method}\}$ with respect to a task y as the absolute difference in frequencies of x co-occurring with y in the same sentence, between years t_1 and t_2 normalized by the total number of entities of the same type as x that co-occur with y in both years. The entity change value is given by: $\delta_y(x)_{t_1}^{t_2} = \frac{|C_{t_1}(x,y) - C_{t_2}(x,y)|}{\sum_{\forall e: \tau(e) = \tau(x)} (C_{t_1}(e,y) + C_{t_2}(e,y))}$, where the frequency of x co-occurring with y in year t is given by $C_t(x, y)$.

In summary, we quantify task trends and research context changes using *task frequency change* and *task stability values*. Below we explore the relationship between *entity change values* and these two variables and estimate the causal impact of TDMM entities on task research landscapes.

5.2 Causal Algorithms

Causal Structure Discovery To uncover the causal structure among variables from observational data, we employ DirectLinGAM ([Shimizu et al., 2011](#)), which assumes a non-Gaussian data-generating process. Since the variables in Section 5.1 come from non-Gaussian frequency distributions, DirectLinGAM is suitable. It uses an entropy-based measure to subtract the effect of each independent variable successively. Unlike PC-Stable ([Colombo and Maathuis, 2014](#)), it does not require iterative search or algorithmic parameters. We apply DirectLinGAM with a 5% significance level for causal discovery (see Appendix C for details).

Causal Inference Once the causal structure between the variables has been established, we leverage this structure to assess the causal effects. Specifically, we measure the causal effects by the *entity change value* of entity x on the *frequency shift* and subsequently on the *stability values* associated with a given task y . For this purpose, we use the probability density function instead of probability mass, as all our causal variables are continuous in nature. We measure the causal effects in two steps: first, we estimate the probability density of the *entity change variable* using a linear regression model. In the next step, we regress the *frequency shift* and *stability* against the *entity change value*, weighted by the inverse probability densities obtained in the previous step. We model the functional form of this regression using a spline to avoid bias due to misspecification. Finally, we calculate the causal effect as [Veitch and Zaveri \(2020\)](#): $\mu(\Delta freq_{t_1}^{t_2}(y)) = \mathbb{E}[\Delta freq_{t_1}^{t_2}(y) | \delta_y(x)_{t_1}^{t_2}]$ and similarly, $\mu(\Delta stability_{t_1}^{t_2}(y)) = \mathbb{E}[\Delta stability_{t_1}^{t_2}(y) | \delta_y(x)_{t_1}^{t_2}]$.

6 Results and Analysis

Correlation-based measures provide a simple way to quantify the association between variables. However, they fall short of explaining complex cause-effect relationships and can yield misleading results. Causality is essential for gaining a deeper understanding of variable relationships, enhancing the robustness and reliability of our findings beyond the limitations of correlation. We discuss more about the importance of causal methods over correlation-based measures in Section 7. In this



Figure 3: Causal Graph of TDMM Entities (entity change values) and Task Entity Frequency Shift.

section, our focus is on uncovering relationships among causal variables (Section 6.1) and measuring the impact of TDMM entities on target task entities (Section 6.2).

6.1 Causal Relation between the Variables

Figure 3 shows the discovered causal graph for the frequency shift of task entities. Overall, we observe that the entity change values of associated tasks, datasets, metrics, and methods have a direct causal effect on the frequency shift values of the target tasks. Since frequency shift value quantifies the trend in NLP research, we infer from the causal graph that the trend of a task is governed primarily by the life cycles of its associated TDMM entities. We see similar causal relation on task stability value (see Figure 4, Appendix A). **Evaluation:** We perform a sensitivity analysis of the causal graph by adding Gaussian noise with zero mean and unit variance to the entity change values in the data (Cinelli et al., 2019). This gives an estimate of the robustness of the graph in the presence of unobserved confounders. We observe that the graph is stable to unobserved confounding, giving all edge probabilities greater than 0.5.

6.2 Causal Impact of the Variables

The organizers of ACL 2018⁴ categorize NLP research into 21 areas, and provide a set of popular tasks for each area. Out of those, we curate 16 areas and select one task from each based on its frequency of occurrence in our corpus. We estimate the effect of TDMM entities (entity change value) behind the development of these tasks (frequency shift value) (see Section 5.1) and summarize the results in Table 4. Since we do not have confounders (Section 6.1), evaluating the causal effect reduces to estimating the conditional expectation of the frequency shift values given the entity change values. We present detailed results in Appendix A.2. We examine the results by addressing the following set of inquiries.

Q1. What role do the methodologies play in causally driving the shift in NLP tasks?

New methodologies have a significant influence on research in various areas of Natural Language Processing (NLP). In the field of Language Modeling, we observe a shift in influence between different methodologies over time.

Between 2003 and 2017, Recurrent Neural Networks (RNNs) had the most decisive impact on **Language Modeling** research. However, this trend shifted with the emergence of Transformers, which have since become the dominant influence in research on this task.

Dialogue Systems, which involve automatic response generation, are closely related to Language Modeling. Therefore, research in this area is highly influenced by Generative Models. From 1990 to 2002, Probabilistic Models played a crucial role in shaping Dialogue Systems research, while RNNs took the lead between 2003 and 2017.

Machine Translation, another task related to Language Modeling, requires the generation of the translated text. Naturally, we observe the influence of similar entities in Machine Translation research. Probabilistic Models had the most decisive impact between 1990 and 2002. In recent years (2018-2022), Transformers have emerged as the dominant influence in this research area.

In the field of **Speech Recognition**, Hidden Markov Models (HMMs) have shown a significant influence. HMMs have played a crucial role in shaping Speech Recognition research between 1979 to 2002.

Named Entity Recognition (NER) has also been influenced by Hidden Markov Models, particularly in its early days (1990-2002), as NER is often formulated as a sequence tagging problem. Various parser algorithms were employed to solve the problem in the period between 2003 and 2017.

For **Semantic Parsing**, parser algorithms have been instrumental and have had a significant impact on research in this area. Between 1979 and 1989, Grammar Induction techniques were used to elicit the underlying semantic parse trees.

From 1990 to 2002, researchers employed various statistical models in **Morphological Analysis**, which is evident from our results.

In **Semantic Role Labeling**, Support Vector Machines and Neural Network Models have been widely used to solve this task.

In **Co-reference Resolution**, Neural Network

⁴<https://acl2018.org/call-for-papers/>

Task	Primary Cause				
	1979-1989	1990-2002	2003-2017	2018-2022	1979-2022
Language Modeling	-	-	Recurrent Neural Networks ^M	Transformers ^M	Transformers ^M
Dialogue System	-	Probabilistic Generative Models ^M	Recurrent Neural Networks ^M	MultiWoz ^D	MultiWoz ^D
Machine Translation	-	Probabilistic Generative Models ^M	WMT Data ^D	Transformers ^M	Transformers ^M
Speech Recognition	Hidden Markov Models ^M	Hidden Markov Models ^M	Machine Translation ^T	Machine Translation ^T	Hidden Markov Models ^M
Named Entity Recognition	-	Hidden Markov Models ^M	POS Tagging ^T	Relation Extraction ^T	POS Tagging ^T
POS Tagging	-	Text Classification ^T	Parser Algorithms ^M	Word Segmentation ^T	Word Segmentation ^T
Semantic Parsing	Grammar Induction ^M	Parser Algorithms ^M	Parser Algorithms ^M	Dependency Parsing ^T	Parser Algorithms ^M
Morphological Analysis	-	Statistical Models ^M	Dependency Parsing ^T	UD Treebank ^D	Statistical Models ^M
Semantic Role Labeling	-	-	Support Vector Machines ^M	Neural Network Models ^M	Support Vector Machines ^M
Co-reference Resolution	-	MUC-VI Text Collection ^D	Integer Linear Programming ^M	Neural Network Models ^M	Neural Network Models ^M
Word Sense Disambiguation	-	Wordnet ^D	Maximum Entropy Models ^M	Neural Network Models ^M	Wordnet ^D
Sentiment Analysis	-	-	Twitter Dataset ^D	Text Classification ^T	Text Classification ^T
Argument Mining	-	-	Text Classification ^T	Sentiment Analysis ^T	Sentiment Analysis ^T
Question Answering	Parsing Algorithms ^M	Information Extraction ^T	Information Extraction ^T	Pre-Trained LLMs ^M	Information Extraction ^T
Textual Entailment	-	-	Statistical Models ^M	Pre-Trained LLMs ^M	Pre-Trained LLMs ^M
Summarization	-	Wordnet ^D	Sentence Compression ^T	Pre-Trained LLMs ^M	Pre-Trained LLMs ^M

Table 4: Causal analysis identifies the main drivers (Methods, Tasks, Datasets) of frequency shifts in NLP tasks across four periods, with "-" indicating insufficient data for analysis.

models have gained prominence starting in 2018. However, from 2003 to 2017, Integer Linear Programming was also utilized to address this problem.

Pre-trained Language Models (LLMs) have demonstrated superior performance in several NLP tasks, including **Question Answering**. Researchers have also explored parsing algorithms to parse questions and align them with potential answers.

Furthermore, **Textual Entailment** and **Summarization** have been heavily influenced by pre-trained LLMs between 2018 and 2022, as evident from our results.

Q2. How have changes in data availability contributed to the NLP Research Tasks?

High-quality datasets play a crucial role in advancing NLP research. While new methodologies are important, they cannot fully propel the field forward without the support of high-quality datasets. Researchers understand the significance of dataset quality and actively curate datasets to drive advancements in the field. Our findings further confirm the prevalence of this trend, highlighting the strong emphasis on dataset quality in NLP research.

In the early stages of deep neural models, such as Recurrent Neural Networks (RNNs), the creation of large datasets became essential for efficient model training. Between 2018 and 2022, several datasets were curated, with MultiWoz being the most widely used dataset for research in **Dialogue Systems**.

In the domain of **Machine Translation**, the significance of datasets in shaping research direction cannot be overlooked. The influence of WMT datasets on Machine Translation research is evident from our findings.

For **Morphological Analysis**, the Universal De-

pendency Treebank dataset is frequently used as a benchmark, indicating its importance in driving research in this area.

During the period of 1990-2002, the creation of the MUC-VI dataset played a crucial role in advancing research in **Co-reference resolution**.

In the field of **Sentiment Analysis**, the Twitter dataset holds significant importance in driving research in this domain.

Overall, our analysis underscores the vital role of datasets in shaping and driving research across various NLP tasks.

Q3. Do evaluation metrics drive paradigm shifts in NLP research?

Most NLP tasks rely on a standard set of metrics borrowed from other domains, such as machine learning and computer vision, to evaluate system performance. However, there is limited research dedicated to improving these metrics within the field of NLP, as it often requires theoretical knowledge beyond the scope of NLP itself. Despite this, our analysis in Table 5 reveals some noteworthy exceptions. Metrics explicitly designed for evaluating NLP tasks, such as BLEU and METEOR, have demonstrated significant impact in advancing **Machine Translation** research. Similarly, the metric ROUGE has influenced research in the field of **Summarization**. While perplexity scores are commonly used to measure the generalization capabilities of probability distributions, they are predominantly utilized for evaluating language models in NLP tasks.

Q4. What is the causal impact of cross-pollination of ideas between related NLP tasks?

We consistently observe a pattern of related NLP tasks evolving in tandem, borrowing ideas and tech-

niques from one another. This trend is clearly reflected in our findings. For instance, **Speech Recognition** and **Machine Translation** are linked as researchers explore end-to-end systems that translate speech, and our results show that Machine Translation has had the greatest influence on Speech Recognition research between 2003 and 2022.

Named Entity Recognition (NER) is commonly approached as a sequence tagging problem, and it is influenced by related tasks such as **POS Tagging** (2003-2017) and **Relation Extraction** (2018-2022), as these problems are often jointly solved. Similarly, POS Tagging initially posed as a text classification problem (1990-2002), is significantly impacted by the **word segmentation task**, as evident from our results in the period of 2018-2022.

In recent years (2018-2022), **dependency** and **semantic parsing** have been jointly solved using the same neural model, highlighting the influence of dependency parsing on research in semantic parsing. **Sentiment Analysis** has garnered considerable research interest and is commonly framed as a text classification problem. Additionally, **Argument Mining**, which involves understanding the sentiments behind arguments, is influenced by sentiment analysis. Furthermore, the classification of various argument components, such as claims and evidence, is often approached as text classification problems, as evidenced by our results.

7 Discussion: Correlation and Causation

“correlation does not imply causation”

– Pearson (1892)

Causation and correlation, although related, are distinct concepts. While they can coexist, correlation does not simply imply causation. Causation signifies a direct cause-and-effect relationship, where one action leads to a specific outcome. In contrast, correlation simply indicates that two actions are related in some way, without one necessarily causing the other.

In our work, we focus on causal inference from data. While correlation-based measures provide a straightforward method for quantifying associations between variables, they often fall short when it comes to explaining complex cause-and-effect relationships.

To demonstrate the effectiveness of our framework, we establish a simple baseline using a PMI-based correlation measure (Bouma, 2009). For this analysis, we select **Machine Translation** as our tar-

get task entity due to its prominent presence in our corpus and the NLP research landscape. We calculate the PMI scores of **Machine Translation** with all other TDMM entities. The PMI score represents the probabilities of co-occurrence between two entities in sentences from research papers, normalized by their individual occurrence probabilities.

Interestingly, we find that **accuracy**, an entity of type **metric**, has the highest PMI score with Machine Translation among all other entities. However, it is important to note that accuracy is a widely used metric across various NLP tasks, and it is not specifically developed for machine translation, nor has machine translation influenced the concept of accuracy. This observation emphasizes the insufficiency of relying solely on correlation-based metrics to understand and analyze research influence on an entity.

We observe that relying solely on correlations can lead to misleading results and interpretations. Therefore, in order to understand the influence of associated TDMM entities on NLP Task entities, we utilize causal algorithms that enable us to gain insights into the cause-and-effect dynamics among the variables we study.

8 Concluding Remarks

In this paper, we retrospectively study NLP research from a causal perspective, quantifying research trends of task entities and proposing a systematic framework using causal algorithms to identify key reasons behind the emergence or disappearance of NLP tasks. Our analysis reveals that tasks and methods are the primary drivers of research in NLP, with datasets following their influence, while metrics have minimal impact. It is important to note that in our analysis, we have structured the reference time into four distinct intervals (see Table 1); however, it can be applied to diverse timeframes, ranging from longer periods to brief intervals, including single years. This adaptability, in the context of rapid recent advancements in NLP, allows to zoom in on local trends and developments that might otherwise go unnoticed (such as the influence of in-context learning on NLP tasks).

We believe our causal analysis enhances understanding of the interplay of research entities in NLP, contributing to the growing body of work on causality and NLP (Feder et al., 2021). We provide with additional analysis and insights in Appendix B.

Limitations

This work is centered on NLP research papers from ACL Anthology, with a focus on papers from the “ACL Events” category. The “ACL Events” category encompasses major conferences, workshops, and journals, including ACL, NAACL, EMNLP, EACL, AACL, CL, and TACL. We also include papers published at COLING from the “non-ACL Events” category. Nevertheless, it is important to acknowledge the presence of NLP papers beyond ACL Anthology in AI journals, regional conferences, and preprint servers. Furthermore, we recognize that certain NLP papers may become available on preprint servers before their official publication in peer-reviewed venues. In this study, we focus on ACL Anthology, which can introduce a time lag when assessing the early impact of influential papers released as preprints (e.g., BERT) or only on preprint servers (e.g., RoBERTa). To address such challenges, we leave the curation and inclusion of NLP research papers from these alternative sources for future works.

Our framework requires research papers tagged with entities as input. Hence, the quality of the tags plays a crucial role in the causal inference of our proposed method. The taggers generate noisy outputs and, thus, might require human intervention to denoise the tags. Moreover, causal algorithms require a large amount of data to produce statistically significant results. Hence, research areas that are less explored or newly emerging may not always be suitable for this framework to be applied. Additionally, we highlight that in this work, we do not consider extra-linguistic factors like author affiliations, funding, gender, etc. We leave them for future research work.

Ethics Statement

In this work, we use publicly available data from ACL Anthology and do not involve any personal data. It is important to recognize that, while our framework is data-driven, individual perspectives toward research are inherently subjective. Decisions involving science should consider data as well as ethical, social, and other qualitative factors. Furthermore, we underscore that the low influence of TDMM entities in our analysis should not be the sole reason for devaluing research papers or reducing their investments. Ethical and academic considerations should guide decisions on research evaluation and resource allocation.

Acknowledgements

We thank Iliia Kuznetsov for his feedback on the initial version of this work. We appreciate all the anonymous reviewers for their helpful comments and suggestions for further analysis. This work has been funded by the German Research Foundation (DFG) as part of the Research Training Group KRITIS No. GRK 2222.

References

- 2008–2022. *Grobid*. <https://github.com/kermitt2/grobid>.
- Alan Akbik, Duncan Blythe, and Roland Vollgraf. 2018. *Contextual string embeddings for sequence labeling*. In *Proceedings of the 27th International Conference on Computational Linguistics*, pages 1638–1649, Santa Fe, New Mexico, USA. Association for Computational Linguistics.
- John M Barrios and Yael Hochberg. 2020. *Risk perception through the lens of politics in the time of the covid-19 pandemic*. Working Paper 27008, National Bureau of Economic Research.
- Antoine Bosselut, Hannah Rashkin, Maarten Sap, Chaitanya Malaviya, Asli Celikyilmaz, and Yejin Choi. 2019. *COMET: Commonsense transformers for automatic knowledge graph construction*. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 4762–4779, Florence, Italy. Association for Computational Linguistics.
- Gerlof Bouma. 2009. Normalized (pointwise) mutual information in collocation extraction. *Proceedings of GSCS*, 30:31–40.
- Kevin W. Boyack and Richard Klavans. 2022. *An improved practical approach to forecasting exceptional growth in research*. *Quantitative Science Studies*, 3(3):672–693.
- Carlos Cinelli, Daniel Kumor, Bryant Chen, Judea Pearl, and Elias Bareinboim. 2019. *Sensitivity analysis of linear structural causal models*. In *Proceedings of the 36th International Conference on Machine Learning*, volume 97 of *Proceedings of Machine Learning Research*, pages 1252–1261. PMLR.
- Diego Colombo and Marloes H. Maathuis. 2014. *Order-independent constraint-based causal structure learning*. *Journal of Machine Learning Research*, 15(116):3921–3962.
- Quang Do, Yee Seng Chan, and Dan Roth. 2011. *Minimally supervised event causality identification*. In *Proceedings of the 2011 Conference on Empirical Methods in Natural Language Processing*, pages 294–303, Edinburgh, Scotland, UK. Association for Computational Linguistics.

- Jesse Dunietz, Lori Levin, and Jaime Carbonell. 2017. [The BECauSE corpus 2.0: Annotating causality and overlapping relations](#). In *Proceedings of the 11th Linguistic Annotation Workshop*, pages 95–104, Valencia, Spain. Association for Computational Linguistics.
- Amir Feder, Katherine A. Keith, Emaad Manzoor, Reid Pryzant, Dhanya Sridhar, Zach Wood-Doughty, Jacob Eisenstein, Justin Grimmer, Roi Reichart, Margaret E. Roberts, Brandon M. Stewart, Victor Veitch, and Diyi Yang. 2021. [Causal inference in natural language processing: Estimation, prediction, interpretation and beyond](#). *CoRR*, abs/2109.00725.
- Christian Fong and Justin Grimmer. 2016. [Discovery of treatments from text corpora](#). In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1600–1609, Berlin, Germany. Association for Computational Linguistics.
- Jonathan Grudin. 2009. [AI and HCI: Two fields divided by a common focus](#). *AI Magazine*, 30(4):48–48.
- David Hall, Daniel Jurafsky, and Christopher D. Manning. 2008. [Studying the history of ideas using topic models](#). In *Proceedings of the 2008 Conference on Empirical Methods in Natural Language Processing*, pages 363–371, Honolulu, Hawaii. Association for Computational Linguistics.
- Yufang Hou, Charles Jochim, Martin Gleize, Francesca Bonin, and Debasis Ganguly. 2019. [Identification of tasks, datasets, evaluation metrics, and numeric scores for scientific leaderboards construction](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 5203–5213, Florence, Italy. Association for Computational Linguistics.
- Yufang Hou, Charles Jochim, Martin Gleize, Francesca Bonin, and Debasis Ganguly. 2021. [TDMSci: A specialized corpus for scientific literature entity tagging of tasks datasets and metrics](#). In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, pages 707–714, Online. Association for Computational Linguistics.
- Zhijing Jin, Zeyu Peng, Tejas Vaidhya, Bernhard Schoelkopf, and Rada Mihalcea. 2021. [Mining the cause of political decision-making from social media: A case study of COVID-19 policies across the US states](#). In *Findings of the Association for Computational Linguistics: EMNLP 2021*, pages 288–301, Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Katherine Keith, David Jensen, and Brendan O’Connor. 2020. [Text and causal inference: A review of using text to remove confounding from causal estimates](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 5332–5344, Online. Association for Computational Linguistics.
- Bernard Koch, Emily Denton, Alex Hanna, and Jacob Gates Foster. 2021. [Reduced, reused and recycled: The life of a dataset in machine learning research](#). In *Thirty-fifth Conference on Neural Information Processing Systems Datasets and Benchmarks Track (Round 2)*.
- Thomas S Kuhn. 1970. *The structure of scientific revolutions*, volume 111. Chicago University of Chicago Press.
- Shixia Liu, Yang Chen, Hao Wei, J. Yang, Kun Zhou, and Steven Mark Drucker. 2015. [Exploring topical lead-lag across corpora](#). *IEEE Transactions on Knowledge and Data Engineering*, 27:115–129.
- Yi Luan, Luheng He, Mari Ostendorf, and Hannaneh Hajishirzi. 2018. [Multi-task identification of entities, relations, and coreference for scientific knowledge graph construction](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 3219–3232, Brussels, Belgium. Association for Computational Linguistics.
- Tomas Mikolov, Ilya Sutskever, Kai Chen, Greg S Corrado, and Jeff Dean. 2013. [Distributed representations of words and phrases and their compositionality](#). In *Advances in Neural Information Processing Systems*, volume 26. Curran Associates, Inc.
- Paramita Mirza and Sara Tonelli. 2014. [An analysis of causality between events and its relation to temporal information](#). In *Proceedings of COLING 2014, the 25th International Conference on Computational Linguistics: Technical Papers*, pages 2097–2106, Dublin, Ireland. Dublin City University and Association for Computational Linguistics.
- Saif M. Mohammad. 2019. [The state of NLP literature: A diachronic analysis of the acl anthology](#). *ArXiv*, abs/1911.03562.
- Saif M. Mohammad. 2020. [Examining citations of natural language processing literature](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 5199–5209, Online. Association for Computational Linguistics.
- Matteo Negri, Marco Turchi, José G. C. de Souza, and Daniele Falavigna. 2014. [Quality estimation for automatic speech recognition](#). In *Proceedings of COLING 2014, the 25th International Conference on Computational Linguistics: Technical Papers*, pages 1813–1823, Dublin, Ireland. Dublin City University and Association for Computational Linguistics.
- Jong-Hoon Oh, Kentaro Torisawa, Chikara Hashimoto, Ryu Iida, Masahiro Tanaka, and Julien Kloetzer. 2016. [A semi-supervised learning approach to why-question answering](#). In *Proceedings of the Thirtieth AAAI Conference on Artificial Intelligence, AAAI’16*, page 3022–3029. AAAI Press.
- Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. [BLEU: a method for automatic evaluation of machine translation](#). In *Proceedings of the*

- 40th annual meeting of the Association for Computational Linguistics, pages 311–318.
- Judea Pearl, Madelyn Glymour, and Nicholas P Jewell. 2016. *Causal inference in statistics: A primer*. John Wiley & Sons.
- Karl Pearson. 1892. The grammar of science. *Nature*, 46(1185):247–247.
- Vinodkumar Prabhakaran, William L. Hamilton, Dan McFarland, and Dan Jurafsky. 2016. [Predicting the rise and fall of scientific topics from trends in their rhetorical framing](#). In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1170–1180, Berlin, Germany. Association for Computational Linguistics.
- Maarten Sap, Ronan Le Bras, Emily Allaway, Chandra Bhagavatula, Nicholas Lourie, Hannah Rashkin, Brendan Roof, Noah A. Smith, and Yejin Choi. 2019. [Atomic: An atlas of machine commonsense for if-then reasoning](#). *Proceedings of the AAAI Conference on Artificial Intelligence*, 33(01):3027–3035.
- Isabel Segura-Bedmar, Paloma Martínez, and María Herrero-Zazo. 2013. [SemEval-2013 task 9 : Extraction of drug-drug interactions from biomedical texts \(DDIExtraction 2013\)](#). In *Second Joint Conference on Lexical and Computational Semantics (*SEM), Volume 2: Proceedings of the Seventh International Workshop on Semantic Evaluation (SemEval 2013)*, pages 341–350, Atlanta, Georgia, USA. Association for Computational Linguistics.
- Naoki Shibata, Yuya Kajikawa, Yoshiyuki Takeda, and Katsumori Matsushima. 2008. [Detecting emerging research fronts based on topological measures in citation networks of scientific publications](#). *Technovation*, 28(11):758–775.
- Shohei Shimizu, Takanori Inazumi, Yasuhiro Sogawa, Aapo Hyvarinen, Yoshinobu Kawahara, Takashi Washio, Patrik O Hoyer, Kenneth Bollen, and Patrik Hoyer. 2011. Directlingam: A direct method for learning a linear non-gaussian structural equation model. *Journal of Machine Learning Research-JMLR*, 12(Apr):1225–1248.
- Janvijay Singh, Mukund Rungta, Diyi Yang, and Saif Mohammad. 2023. [Forgotten knowledge: Examining the citational amnesia in NLP](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 6192–6208, Toronto, Canada. Association for Computational Linguistics.
- Henry Small. 2006. Tracking and predicting growth areas in science. *Scientometrics*, 68(3):595–610.
- Chenhao Tan, Dallas Card, and Noah A. Smith. 2017. [Friendships, rivalries, and trysts: Characterizing relations between ideas in texts](#). In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 773–783, Vancouver, Canada. Association for Computational Linguistics.
- Ana Sabina Uban, Cornelia Caragea, and Liviu P. Dinu. 2021. [Studying the evolution of scientific topics and their relationships](#). In *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021*, pages 1908–1922, Online. Association for Computational Linguistics.
- 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*, volume 30. Curran Associates, Inc.
- Victor Veitch, Dhanya Sridhar, and David Blei. 2020. [Adapting text embeddings for causal inference](#). In *Proceedings of the 36th Conference on Uncertainty in Artificial Intelligence (UAI)*, volume 124 of *Proceedings of Machine Learning Research*, pages 919–928. PMLR.
- Victor Veitch and Anisha Zaveri. 2020. [Sense and sensitivity analysis: Simple post-hoc analysis of bias due to unobserved confounding](#). In *Advances in Neural Information Processing Systems*, volume 33, pages 10999–11009. Curran Associates, Inc.
- Jan Philip Wahle, Terry Ruas, Mohamed Abdalla, Bela Gipp, and Saif M. Mohammad. 2023. [We are who we cite: Bridges of influence between natural language processing and other academic fields](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, Singapore, Singapore. Association for Computational Linguistics.
- Laura Wendlandt, Jonathan K. Kummerfeld, and Rada Mihalcea. 2018. [Factors influencing the surprising instability of word embeddings](#). 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 2092–2102, New Orleans, Louisiana. Association for Computational Linguistics.
- Christopher KI Williams and Carl Edward Rasmussen. 2006. *Gaussian processes for machine learning*, volume 2. MIT press Cambridge, MA.
- Dongxiang Zhang, Lei Wang, Nuo Xu, Bing Tian Dai, and Heng Tao Shen. 2018. [The gap of semantic parsing: A survey on automatic math word problem solvers](#). *CoRR*, abs/1808.07290.

A Appendix: Additional Results

A.1 Causal Relation

In Figure 4, we observe that entity change values of tasks, datasets, metrics and methods have direct causal influence on task stability value.




Figure 4: **Causal Graph:** The graph shows that the emergence and disappearance of TDMM entities (entity change values) have a direct causal effect on the stability of task entities.

A.2 Causal Effects

In Table 5, we observe the entities (Tasks, Datasets, Methods, and Metrics) that influence research on a given NLP Task.

B Appendix: Supplementary Analysis

In addition to the primary results presented in the paper (Section 6), in this section, we describe the supplementary analysis.

B.1 NLP Tasks and Their Dataset Evolution

Frequently Pursued NLP Tasks. From Table 5 in our paper, we observe that overall (from 1979-2022), among all the tasks, “Text Classification” (column 6) holds a remarkable position. This prominence stems from the frequent usage of various NLP tasks being framed or aligned as “Text Classification” or borrowing concepts from it to address other tasks such as “Sentiment Analysis” or “Word Sense Disambiguation.” Additionally, our framework offers the flexibility to perform a similar analysis between any chosen periods.

Evolution of Datasets in NLP Tasks. Referring to Table 5 in our paper, in the context of “Speech Recognition,” we observe a shift in influential datasets over different periods. Between 1990-2002, the “WSJ Corpus” took the lead, while in the subsequent period of 2003-2017, the “ATIS Dataset” had more influence. Interestingly, between 2018-2022, the trend shifted once again to the “Switchboard Dataset”.

A similar trend is reflected in the “Summarization” task as well: in the years 1990-2002, “Wordnet” played a significant role, while the “Gigaword Dataset” took over in 2003-2017. However, in the most recent period of 2018-2022, “Pubmed” emerged as the notable dataset for the “Summarization” task.

Common Datasets Across NLP Tasks. We observe from Table 5 (column 6) that across the entire span from 1979 to 2022, the “Penn Tree-

bank” dataset emerged as a pivotal influence, significantly impacting tasks such as “Language Modeling,” “POS Tagging,” and “Semantic Parsing.” Using our framework, a similar analysis could also be done between any chosen periods.

B.2 Entity Influence on Task Frequency and Stability

Influence of Research Entities on Task Stability.

We measure the causal effect of research entities on Task Stability Value (see Section 5.1). From the resulting causal graph (Figure 4), we observe that the entity change values of associated tasks, datasets, metrics, and methods directly impact the stability value of the target task, similar to the task frequency shift value.

Correlations Between Task Frequency Change and Stability.

We observe a slightly positive correlation between frequency change and stability of research tasks with a Pearson coefficient of 0.08. This is because when a new task emerges, initially, a few researchers start working on it, which gradually increases its frequency of appearance. At the same time, researchers experiment with various methods and datasets to solve these newly emerged tasks, causing high instability (e.g., Math Problem Solving (Zhang et al., 2018)). On the contrary, the opposite is not always true: well-defined tasks are often the most researched, and yet researchers always explore new ideas on these tasks, which harms stability.

Overview and Insights. Our analysis shows that research in NLP is primarily driven by tasks and methods; the influence of datasets follows them, and metrics have minimum impact. Our analysis of frequency shift values reveals the gradual paradigm shift in NLP research. Initially, the focus was on practical problems such as **Speech Recognition** and **Machine Translation**. However, over time, researchers ventured into more complex areas like **textual entailment** and **argument mining**, necessitating domain knowledge and extensive data rea-

Task	Primary Cause				
	1979-1989	1990-2002	2003-2017	2018-2022	1979-2022
Language Modeling	-	-	Recurrent Neural Networks ^M	Transformers ^M	Transformers ^M
	-	-	Machine Translation ^T	Text Generation ^T	Text Generation ^T
	-	-	Penn Treebank ^D	Perplexity ^m	Perplexity ^m
	-	-	Perplexity ^m	SuperGLUE ^D	Penn Treebank ^D
Dialogue System	-	-	Recurrent Neural Networks ^M	MultiWoz ^D	MultiWoz ^D
	-	-	MultiWoz ^D	Transformers ^M	Transformers ^M
	-	-	Language Generation ^T	Response Generation ^T	Response Generation ^T
	-	-	Perplexity ^m	Rouge ^m	Rouge ^m
Machine Translation	-	Probabilistic Generative Models ^M	WMT Data ^D	Transformers ^M	Transformers ^M
	-	Speech Recognition ^T	BLEU ^m	METEOR ^m	METEOR ^m
	-	Perplexity ^m	Attention Mechanism ^M	Language Modeling ^T	Language Generation ^T
	-	Penn Treebank ^D	Language Generation ^T	WMT Data ^D	WMT Data ^D
Speech Recognition	Hidden Markov Models ^M	Hidden Markov Models ^M	Machine Translation ^T	Machine Translation ^T	Hidden Markov Models ^M
	Machine Translation ^T	WSJ Corpus ^D	Hidden Markov Models ^M	Acoustic Models ^M	Language Modeling ^T
	Perplexity ^m	Perplexity ^m	ATIS Dataset ^D	Switchboard Dataset ^D	Perplexity ^m
	-	Language Modeling ^T	Word Error Rate ^m	Word Error Rate ^m	ATIS Dataset ^D
Named Entity Recognition	-	Hidden Markov Models ^M	POS Tagging ^T	Relation Extraction ^T	POS Tagging ^T
	-	Information Extraction ^T	Conditional Random Fields ^M	Wikipedia Corpus ^D	Random Fields ^M
	-	Genia Corpus ^D	Pubmed ^D	Pre-Trained LLMs ^M	Conditional Ontonotes ^D
	-	F1 Score ^m	F1 Score ^m	F1 Score ^m	F1 Score ^m
POS Tagging	-	Text Classification ^T	Parser Algorithms ^M	Word Segmentation ^T	Word Segmentation ^T
	-	Discriminative Models ^M	Word Segmentation ^T	Neural Network Models ^M	Neural Network Models ^M
	-	Penn Treebank ^D	Penn Treebank ^D	Penn Treebank ^D	Penn Treebank ^D
	-	F1 Score ^m	F1 Score ^m	F1 Score ^m	F1 Score ^m
Word Sense Disambiguation	-	Wordnet ^D	Maximum Entropy Models ^M	Neural Network Models ^M	Wordnet ^D
	-	Semantic Tagging ^T	Text Classification ^T	Text Classification ^T	Neural Network Models ^M
	-	Discriminative Models ^M	Wordnet ^D	Wordnet ^D	Text Classification ^T
	-	Accuracy ^m	F1 Score ^m	F1 Score ^m	F1 Score ^m
Morphological Analysis	-	Statistical Models ^M	Dependency Parsing ^T	UD Treebank ^D	Statistical Models ^M
	-	Word Segmentation ^T	Statistical Models ^M	Pre-Trained LLMs ^M	Dependency Parsing ^T
	-	UD Treebank ^D	UD Treebank ^D	Lemmatization ^T	UD Treebank ^D
	-	Accuracy ^m	Accuracy ^m	F1 Score ^m	Accuracy ^m
Semantic Parsing	Grammar Induction ^M	Parser Algorithms ^M	Parser Algorithms ^M	Dependency Parsing ^T	Parser Algorithms ^M
	Information Retrieval ^T	Information Extraction ^T	Dependency Parsing ^T	Parser Algorithms ^M	Penn Treebank ^D
	Accuracy ^m	Penn Treebank ^D	Penn Treebank ^D	Penn Treebank ^D	Dependency Parsing ^T
	-	F1 Score ^m	F1 Score ^m	F1 Score ^m	F1 Score ^m
Semantic Role Labeling	-	-	Support Vector Machines ^M	Neural Network Models ^M	Support Vector Machines ^M
	-	-	Relation Extraction ^T	Named Entity Recognition ^T	Named Entity Recognition ^T
	-	-	Propbank ^D	Propbank ^D	Propbank ^D
	-	-	F1 Score ^m	F1 Score ^m	F1 Score ^m
Co-reference Resolution	-	MUC-VI Text Collection ^D	Integer Linear Programming ^M	Neural Network Models ^M	Neural Network Models ^M
	-	Discriminator Models ^M	Ontonotes ^D	Ontonotes ^D	Ontonotes ^D
	-	Word Sense Disambiguation ^T	Mention Detection ^T	Mention Detection ^T	F1 Score ^m
	-	F1 Score ^m	F1 Score ^m	F1 Score ^m	Mention Detection ^T
Sentiment Analysis	-	-	Twitter Dataset ^D	Text Classification ^T	Text Classification ^T
	-	-	Text Classification ^T	Pre-Trained LLMs ^M	Neural Network Models ^M
	-	-	Neural Network Models ^M	Amazon Reviews ^D	Twitter Dataset ^D
	-	-	F1 Score ^m	F1 Score ^m	F1 Score ^m
Argument Mining	-	-	Text Classification ^T	Sentiment Analysis ^T	Sentiment Analysis ^T
	-	-	Neural Network Models ^M	Neural Network Models ^M	Neural Network Models ^M
	-	-	Wikipedia Corpus ^D	Wikipedia Corpus ^D	Wikipedia Corpus ^D
	-	-	F1 Score ^m	F1 Score ^m	F1 Score ^m
Question Answering	Parsing Algorithms ^M	Information Extraction ^T	Information Extraction ^T	Pre-Trained LLMs ^M	Information Extraction ^T
	Information Retrieval ^T	Wordnet ^D	Freebase ^D	Squad ^D	Pre-Trained LLMs ^M
	Accuracy ^m	Accuracy ^m	Parsing Algorithms ^M	Summarization ^T	Squad ^D
	-	Statistical Models ^M	F1 Score ^m	F1 Score ^m	F1 Score ^m
Textual Entailment	-	-	Statistical Models ^M	Pre-Trained LLMs ^M	Pre-Trained LLMs ^M
	-	-	Information Extraction ^T	SNLI Dataset ^D	SNLI Dataset ^D
	-	-	F1 Score ^m	Text Classification ^T	Text Classification ^T
	-	-	-	F1 Score ^m	F1 Score ^m
Summarization	-	Wordnet ^D	Sentence Compression ^T	Pre-Trained LLMs ^M	Pre-Trained LLMs ^M
	-	Probabilistic Generative Models ^M	Recurrent Neural Networks ^M	Rouge ^m	Rouge ^m
	-	F1 Score ^m	Rouge ^m	Pubmed ^D	Question Answering ^T
	-	Information Retrieval ^T	Gigaword ^D	Question Answering ^T	Pubmed ^D

Table 5: **The primary reason behind the frequency shift of the tasks.** We analyze the trends in four different periods of reference. Most influential Task(T), Dataset(D), Method(M) and Metric(m) are given in the decreasing order of their influence. "-" means there is not enough data instances for the causal analysis.

soning. Examining stability values, we note that pre-trained language models have emerged as versatile solutions, reducing the need for task-specific approaches.

C Appendix: Algorithms

C.1 DirectLinGAM

In Algorithm 1, we describe the DirectLinGAM algorithm (oracle version) in high level as described by Shimizu et al. (2011).

D Appendix: Multiple Linear Regression

We use multiple linear regression to regress a variable on several variables (Pearl et al., 2016). For instance, if we want to predict the value of a variable Y using the values of variables $X_1, X_2, \dots, X_{k-1}, X_k$, we perform multiple linear regression of Y on $\{X_1, X_2, \dots, X_{k-1}, X_k\}$, and estimate a regression relationship (Eqn. 1), which represents an inclined plane through the $(k + 1)$ -dimensional coordinate system.

$$Y = r_0 + \sum_{i=1}^k r_i X_i \quad (1)$$

The Gauss-Markov theorem (Williams and Rasmussen, 2006) simplifies the computation of partial regression coefficients (r_1, \dots, r_k in Eqn 1). It states that if we write Y as a linear combination of $X_1, X_2, \dots, X_{k-1}, X_k$ and noise term ϵ ,

$$Y = r_0 + \sum_{i=1}^k r_i X_i + \epsilon \quad (2)$$

then, regardless of the distributions of the variables Y, X_1, X_2, \dots, X_k , the best least-square coefficients are obtained when ϵ is uncorrelated with each regressors, i.e.,

$$Cov(\epsilon, X_i) = 0, \forall i = 1, 2, \dots, k \quad (3)$$

Algorithm 1: Causal Graph Discovery: DirectLinGAM-Algorithm

- 1 Given a p -dimensional random vector x , a set of its variable subscripts U and a $p \times n$ data matrix of the random vector as X , initialize an ordered list of variables $K := \phi$ and $m := 1$;
 - 2 Repeat until $p - 1$ subscripts are appended to K : Perform least square regression of x_i and $x_j, \forall i \in U - K (i \neq j)$ and compute the residual vectors $r^{(j)}$ and the residual data matrix $R^{(j)}$ from the matrix $X, \forall j \in U - K$. Final a variable x_m independent of its residuals and append m to the end of K ;
 - 3 Append the remaining variable to the end of K ;
 - 4 Construct a strictly lower triangular matrix B by following the order in K , and estimate the connection strengths b_{ij} by using some conventional covariance-based regression such as least squares and maximum likelihood approaches on the original random vector x and the original data matrix X ;
-