Artificial Intelligence: Opportunities for Prevention Science

Jonathan Wright
Alex Winters

Topics for Today

Machine learning for social science

How things are currently being done

How AI/ML will change the fields of social science

Bridging the gap between researchers, policy makers, and effective AI/ML applications

Machine Learning for Social Science: An Agnostic Approach (Grimmer et al., 2021)

- Deductive -> Inductive

  Social science is moving from a deductive approach, centered on individual experimentation, to a more sequential, interactive, inductive approach.

Machine Learning for Social Science: An Agnostic Approach (Grimmer et al., 2021)

- Data Scarcity -> Massive Datasets

  This shift is driven by data abundance and the increasing computational power available for analysis.

Machine Learning for Social Science: An Agnostic Approach (Grimmer et al., 2021)

- Data Extraction with ML

  Machine learning methods are used to extract meaning and discover new patterns in these datasets, offering new opportunities for social science.

Machine Learning for Social Science: An Agnostic Approach (Grimmer et al., 2021)

- An Agnostic Approach

  There isn’t one best ML method. We use methods that best suit the task at hand, whether it is discovery, measurement, or causal inference.

Machine Learning for Social Science: An Agnostic Approach (Grimmer et al., 2021)

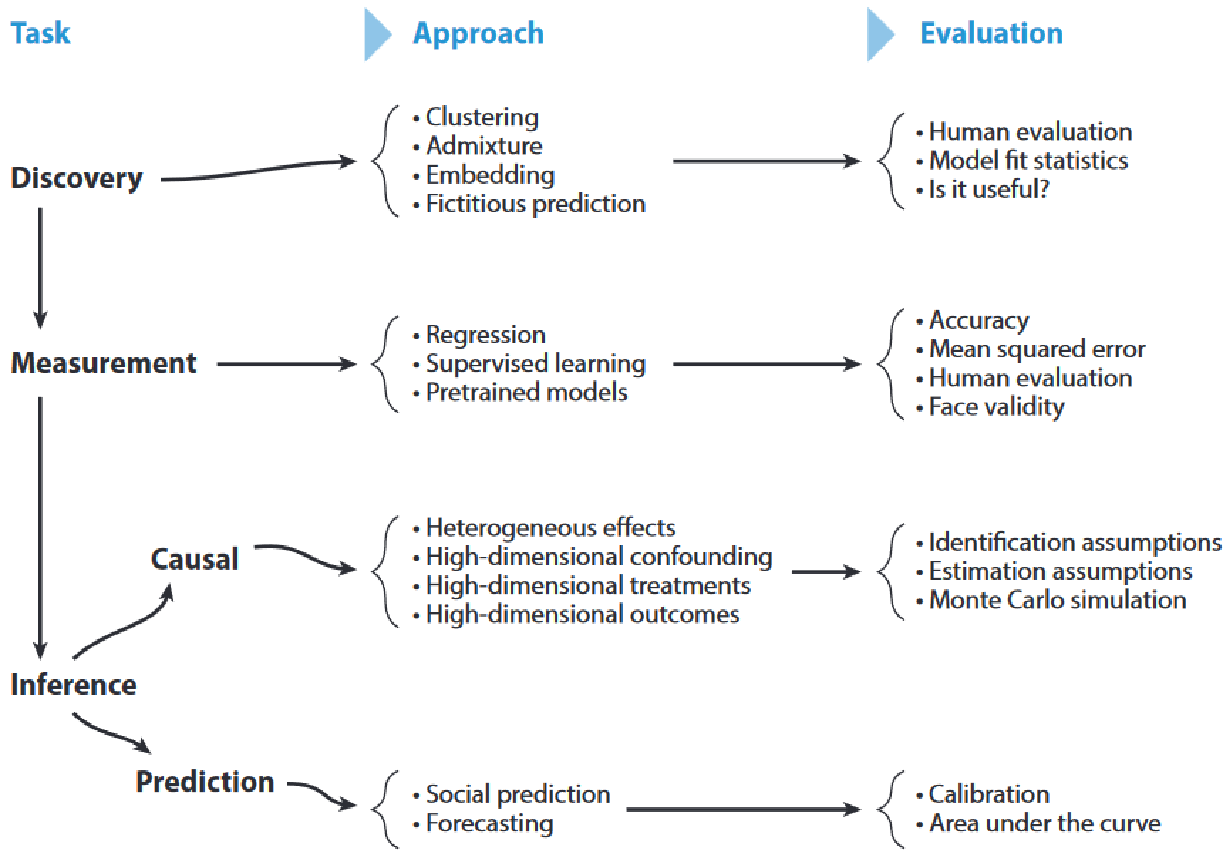

Figure 1

Our approach to machine learning in the social sciences. We reframe the tasks to ones relevant to social science: discovery, measurement, causal inference, and prediction.

Machine Learning for Social Science: An Agnostic Approach (Grimmer et al., 2021)

- Discovery

  Quantitative empirical work tends to ignore the process of concept formation resulting from data interaction. Machine learning procedures expand our tools for engaging in data-driven discovery of new concepts that underly data (e.g., Unsupervised ML, such as clustering, admixture models, and embeddings).

Machine Learning for Social Science: An Agnostic Approach (Grimmer et al., 2021)

- Measurement and Validation

  ML enables custom measurements (e.g., automatic coding, synthetic datasets, etc.). This democratizes research, but new measurement strategies require that researchers provide evidence that their measures capture the theoretical concepts of interest and are accurate.

Machine Learning for Social Science: An Agnostic Approach (Grimmer et al., 2021)

- Causal Inference and Prediction

  Careful consideration must be given to the type of research question because prediction (forecasting outcomes) and causal inference (understanding the effect of an intervention) require different ML methods.

Machine Learning for Social Science: An Agnostic Approach (Grimmer et al., 2021)

- Conclusion

  The abundance of data offers new opportunities for iterative, inductive research methods. Machine learning tools allow social scientists to discover and measure new concepts, making it essential to adapt methods to the task at hand.

ML Methods in Common Parlance

Category	Common Models/Algorithms	Applications in Social Science
Supervised Learning	Linear Regression, Logistic Regression, Decision Trees, Random Forest, SVM, Naive Bayes	Predicting outcomes, classification tasks like voting behavior
Unsupervised Learning	K-Means, Hierarchical Clustering, LDA, PCA, t-SNE	Clustering groups, topic modeling, dimensionality reduction
Semi-Supervised Learning	Self-Training, Graph-Based Learning	Survey analysis, network analysis
NLP Models	Word2Vec, BERT, Text Classification	Analyzing political rhetoric, sentiment analysis, topic discovery
Causal Inference	DiD, PSM, Instrumental Variables	Estimating causal effects in policy studies, economic research
Ensemble Methods	Random Forest, Gradient Boosting, Bagging	Improving prediction accuracy in complex social science datasets
Deep Learning	CNNs, RNNs	Analyzing text, social media data, understanding public opinion trends

Supervised Learning Models

These models are used when the dataset has labeled outcomes (i.e., the target variable is known), and the goal is to predict or classify new observations based on past data.

- **Linear Regression**: Used for predicting continuous outcomes, such as economic indicators or survey scores.
- **Logistic Regression**: Common for binary classification tasks, like predicting political party affiliation or the likelihood of an event (e.g., voting behavior).
- **Decision Trees**: These models break down decisions into a tree-like structure, useful in examining causal relationships (e.g., examining factors that affect policy adoption).
- **Random Forest**: An ensemble method that builds multiple decision trees and averages their results for classification and regression, widely used for its robustness in predicting outcomes.
- **Support Vector Machines (SVM)**: Typically used for classification tasks, such as identifying ideological positions based on text data.
- **Naive Bayes**: Often applied to text classification, including sentiment analysis, topic identification, or classifying political speeches.
- **Gradient Boosting Machines (GBM)**: Like random forests but using a boosting approach to correct errors iteratively, often used in social science for predicting election results or economic outcomes.
- **LASSO (Least Absolute Shrinkage and Selection Operator)**: A form of regression that performs both variable selection and regularization, popular in high-dimensional social science data.

Unsupervised Learning Models

Unsupervised learning involves finding hidden patterns or structures in data where there are no predefined labels or outcomes.

- **K-Means Clustering**: Used for grouping similar observations into clusters, such as grouping survey respondents based on voting patterns or consumer behavior.
- **Hierarchical Clustering**: Similar to K-means but builds a hierarchy of clusters, often used in network analysis or social groupings.
- **Latent Dirichlet Allocation (LDA)**: A widely-used topic modeling algorithm for discovering topics in text data, such as analyzing political speeches or newspaper articles.
- **Principal Component Analysis (PCA)**: A dimensionality reduction technique used to find the underlying factors that explain variation in data, often used in survey analysis and election studies.
- **t-Distributed Stochastic Neighbor Embedding (t-SNE)**: A visualization tool used to explore high-dimensional data, often for visualizing clusters in political or social data.
- **Admixture Models**: Used to identify proportions of membership across multiple latent categories, often applied in cultural and network studies.

Semi-Supervised Learning Models

Semi-supervised learning combines labeled and unlabeled data, which is common in social science where only some data points may have labels.

- **Self-Training**: A semi-supervised method that iteratively assigns labels to the unlabeled data, often used in surveys where only some responses are categorized.
- **Graph-Based Semi-Supervised Learning**: Used in network analysis where some nodes are labeled (e.g., political actors) and others are not, but relationships between them can be leveraged to make predictions.

Natural Language Processing (NLP) Models

These models are increasingly used in social science research due to the abundance of text data (e.g., political speeches, social media posts).

- **Word2Vec / GloVe**: Word embedding models used to represent words in a vector space, allowing for similarity comparisons (e.g., analyzing shifts in political rhetoric).
- **BERT (Bidirectional Encoder Representations from Transformers)**: A transformer-based language model for tasks like sentiment analysis, text classification, and even causal inference from text data.
- **Text Classification Models**: Models like SVM, Naive Bayes, or neural networks are applied to classify text data, including detecting fake news, analyzing sentiment, or identifying political ideology.

Deep Learning Models

These models, although less common in traditional social science, are becoming more popular with larger datasets.

- **Convolutional Neural Networks (CNNs)**: While often used for image data, CNNs can be applied to text data in social sciences for classification tasks.
- **Recurrent Neural Networks (RNNs)**: Typically used for sequence data, such as time series or text analysis, they are increasingly being applied to understand dynamic social phenomena like public opinion trends.

Comparing Models

A common approach to learning and validation includes the comparison of models on the same task.

- **Convolutional Neural Networks (CNNs)**: While often used for image data, CNNs can be applied to text data in social sciences for classification tasks.
- **Naive Bayes**: Often applied to text classification, including sentiment analysis, topic identification, or classifying political speeches.

The Impact of Pause Types on Adverse Listening Condition Classification with Convolutional Neural Networks and Naïve Bayes

Jonathan Wright and Melissa Baese-Berk

Department of Linguistics, University of Oregon

61st Annual Meeting

VIRTUAL PSYCHONOMICS

Background and Research Questions

- Adverse listening conditions, such as speech in noise and non-native directed speech, create difficulty for talkers in conversations, leading to interlocutor focused speech modulations (i.e., clear speech).
- Analyses of recently developed clear speech corpora (e.g., LUCID corpus ; Hazan and Baker, 2011) suggest that talkers modulate phonetic features according to interlocutors' needs.
- Prior research also suggests that talkers modulate the frequency and duration of pauses during clear speech, which may contribute to increased intelligibility (Bradlow et al., 2003; Smiljanić & Bradlow, 2008).
- However, the role of silences and filled pauses, across different types of adverse listening conditions has not been studied.

? Does a Convolutional Neural Net model, which accounts for lexical neighborhoods, more accurately predict the type of adverse listening condition than a Naïve Bayes model, which only accounts for individual word probabilities?

? Does the presence of silences and filled pauses significantly impact adverse listening condition prediction?

Dataset and Methods

- Lucid corpus: 240 conversations: babble noise (n=60), non-native interlocutor (n=60), vocoder noise (n=60), and no barrier (n=60; a control condition). From annotated transcriptions, text from the talker forced to produce clear speech was extracted.
- To perform predictions we extracted “sentences” by splitting the text on the first silence after 100 characters, resulting in 10,935 sentences: babble (n=2639), non-native (n=3476), vocoder (n=2675), no barrier (n=2145). Note: n was not equal across conditions
- Naïve Bayes and 1D CNN were compared across 4 conditions: silences removed, filled pauses (e.g. er, erm, um, mm) removed, all pauses removed, no pauses removed.
- Naïve Bayes: independent lexical items extracted into a “bag of words”. Train: 80%. Test: 20%. Given the occurrence of the words in each condition, probabilities were used to predict the condition.
- 1D Convolutional Neural Network: lexical items are dependent on lexical neighborhoods. Train: 80%. Test: 20%. 1D CNN model: two Conv1D filter layers, each followed by a pooling layer, one hidden layer, one output layer with a node for each type of activity, three epochs, batch size of 64.

Results and Discussion

	Naïve Bayes	1D CNN
Full Corpus	47.2%	59.8%
No Silences	47.5%	57.8%
No Filled Pauses	44.6%	54.6%
No Pauses	45.1%	53.6%

Figure 1: Prediction results for Naïve Bayes and 1D CNN models across four conditions. Chance is at 25%.

Both models accurately predicted adverse listening conditions above chance, suggesting:

- NB: speakers alter lexical choices depending on the adverse listening condition of the interlocutor.
- CNN: greater accuracy suggests speakers alter lexical and syntactic choices.

Pause Conditions:

- Silences: did not benefit NB, some benefit for CNN
- Filled pauses: impacted prediction more than silence in NB and CNN

Results and Discussion

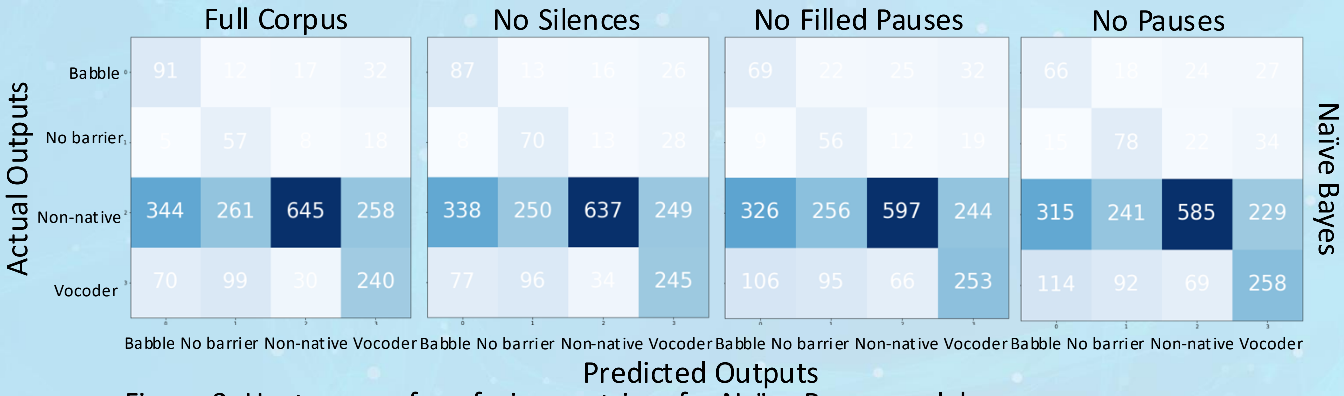

Figure 3: Heat maps of confusion matrices for Naïve Bayes models

	Filled Pauses
Babble	.051%
No barrier	.041%
Non-native	.078%
Vocoder	.041%

Figure 4: Percent of filled pauses in each clear speech type.

- NB most accurately predicted the non-native listening condition, suggesting that lexical items differed the most when speaking to a non-native interlocutor.
- Further, accuracy dropped substantially when filled pauses were removed.
- Are filled pauses especially utilized in non-native directed speech?

- Filled pause percent and confusion matrices correspond. As pauses are removed, accuracy is reduced on babble and non-native conditions and is improved on vocoder and no barrier conditions. This is less pronounced with the CNNs.

Results and Discussion

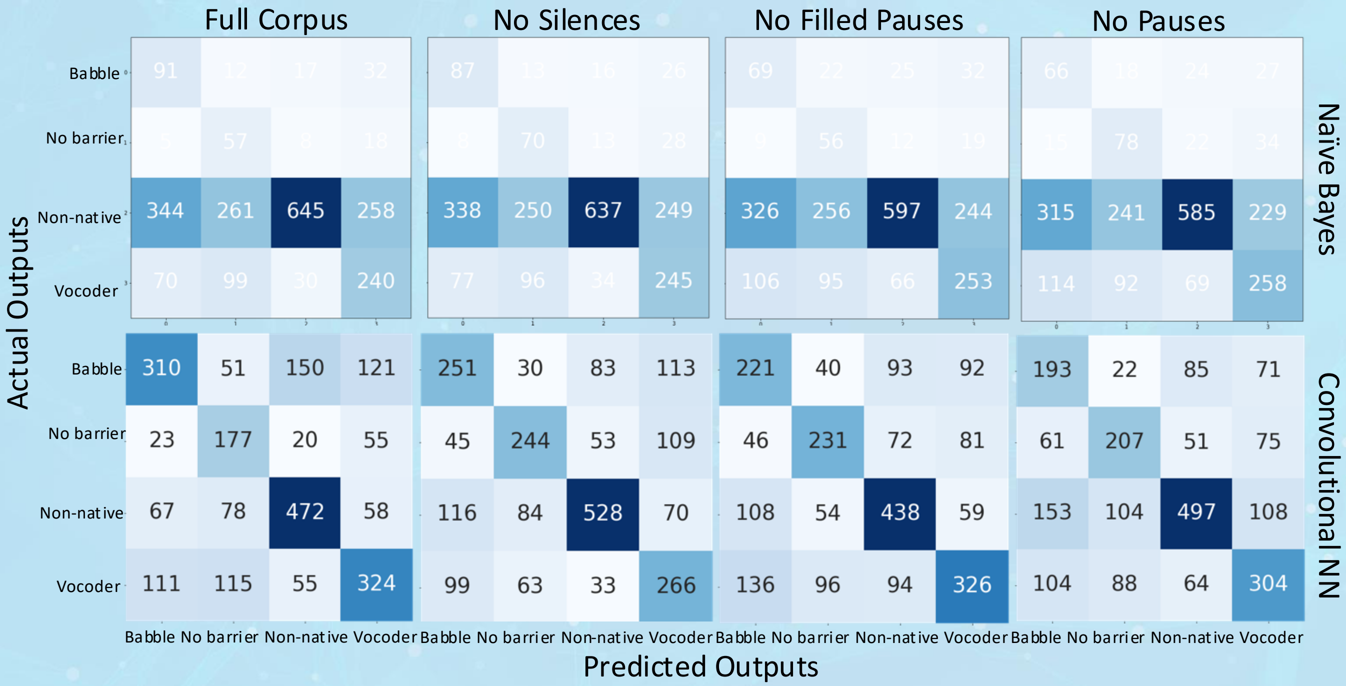

- Like NB, CNN most accurately predicted the non-native listening category. Thus, clear speech with non-native interlocutors may differ substantially from speech in noise.
- While NB defaulted to predicting the non-native category, CNN did not, thereby predicting each category above chance. This may suggest that CNNs do better overall with uneven sample sizes.

(1147) Wright, Baese-Berk

What does this enable for researchers

- Ability to better handle unstructured text data.
- Modeling Long-Term Dependencies and Relationships
- Enhanced Predictive Modeling
- Multimodal Data Integration
- Scalable Processing

How do the soft sciences benefit the most

- Super-charged NLP
  - Immense amount of information that exists within text
- Qualitative to Quantitative
- Immense access and wealth of text-based data
  - Social Media
  - Health Care Data
  - Narrative Social Welfare Agencies
  - Transcribed Interviews
  - Survey Data
- New Avenues
  - Opportunity for analysis and designing experiments and research questions around AI/ML

Transformers Overview

- Transformers
  - Sequence to sequence models
  - Don't require rigid and immutable inputs
  - Embeds into a high dimensional vector space
- Attention
  - A mechanism that allows tokens and inputs to understand the context of the sequence that they exist in
  - Example: Being able to understand that 'bank' has two different meanings in the sentence "The bandit robbed the bank, and then ran all the way down to the bank of the river."

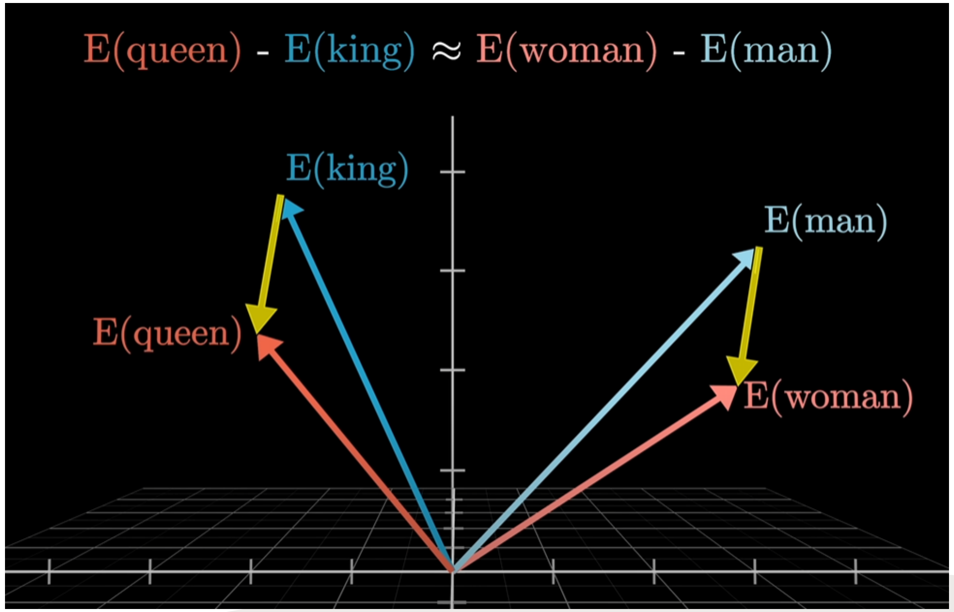

Transparency, Interpretability, Explainability

- Transparency
  - Overall model structure
  - Individual components
  - Learning algorithm
- Interpretations are mappings of abstract concepts, that exist in the model, into a domain that can make sense to a human
- No Concise definition of explainability
  - Explanations can differ in completeness and degree of causality
  - Domain knowledge is a central part to explainability

Bridging the Gap

Creating	Creating more robust systems for prevention researchers to apply AI/ML methods
Developing	Developing institutional architecture and expertise and these domains
Applying	Applying for seed funding for AI infrastructure
Buttressing	Buttressing the work of Social Scientists with the support of ML engineering and Data Science

Cloud

High powered compute infrastructures are a necessity

Big Data

Model Training

Fully Managed Machine Learning Environments

Full workflow control

AWS Sagemaker and Microsoft AI Studio

Pay for what you use

Our Projects

Retrieval Augmented Generation	• Chatbots • Enterprise data assistants
Image and Audio Analysis	• Layout Identification and Text extraction • Document Comparison • Text transcription
Natural Language Processing	• Sentiment Analysis • Interview Summarization and Topic Modeling

References

Bradlow, A. R., Kraus, N., & Hayes, E. (2003). Speaking clearly for children with learning disabilities. Journal of Speech, Language, and Hearing Research.

Grimmer, J., Roberts, M. E., & Stewart, B. M. (2021). Machine Learning for Social Science: An Agnostic Approach. Annual Review of Political Science, 24(Volume 24, 2021), 395–419. https://doi.org/10.1146/annurev-polisci-053119-015921

Hazan, V., & Baker, R. (2011). Acoustic-phonetic characteristics of speech produced with communicative intent to counter adverse listening conditions. The Journal of the Acoustical Society of America, 130(4), 2139–2152.

Smiljanić, R., & Bradlow, A. R. (2008). Temporal organization of English clear and conversational speech. The Journal of the Acoustical Society of America, 124(5), 3171–3182.