date : 2026-08-03,22:20

Reference

[YouTube : NLP in FinTech]

NLP in FinTech Applications - AACL-IJCNLP-2020 Tutorial

introduction

related to two paper

- "NLP in FinTech Application"
- "Fine-grained Financial Opinion Mining "

speaker information

- research interest : applying NLP in the finance & economic
- https://linktr.ee/cjchen

FinTech

The Financial Stability Board (FSB)

- FinTech is technology-enabled innovation in financial services that could result in new business models, applications, processes or products with an associated material effect on the provision of financial services.

> Simple : **Using technology to improve financial services.**

- the fundamental services haven't changed. only the way to provide the service change

Financial Services

- Payments
- Deposits and Lending
- Market Provisioning
- Capital Raising
- Insurance
- Investment Management

Overview of FinTech Application

FinTech Applications

NLP Lab
National Taiwan University

- **Payments**
  - Cashless World
  - Emerging Payment Rails
- **Deposits and Lending**
  - Alternative Lending
  - Shifting Customer Preferences
- **Market Provisioning**
  - Smarter, Faster Machines
  - New Market Platforms
- **Capital Raising**
  - Crowdfunding
- **Insurance**
  - Insurance Disaggregation
  - Connected Insurance
- **Investment Management**
  - Empowered Investors
  - Process Externalisation

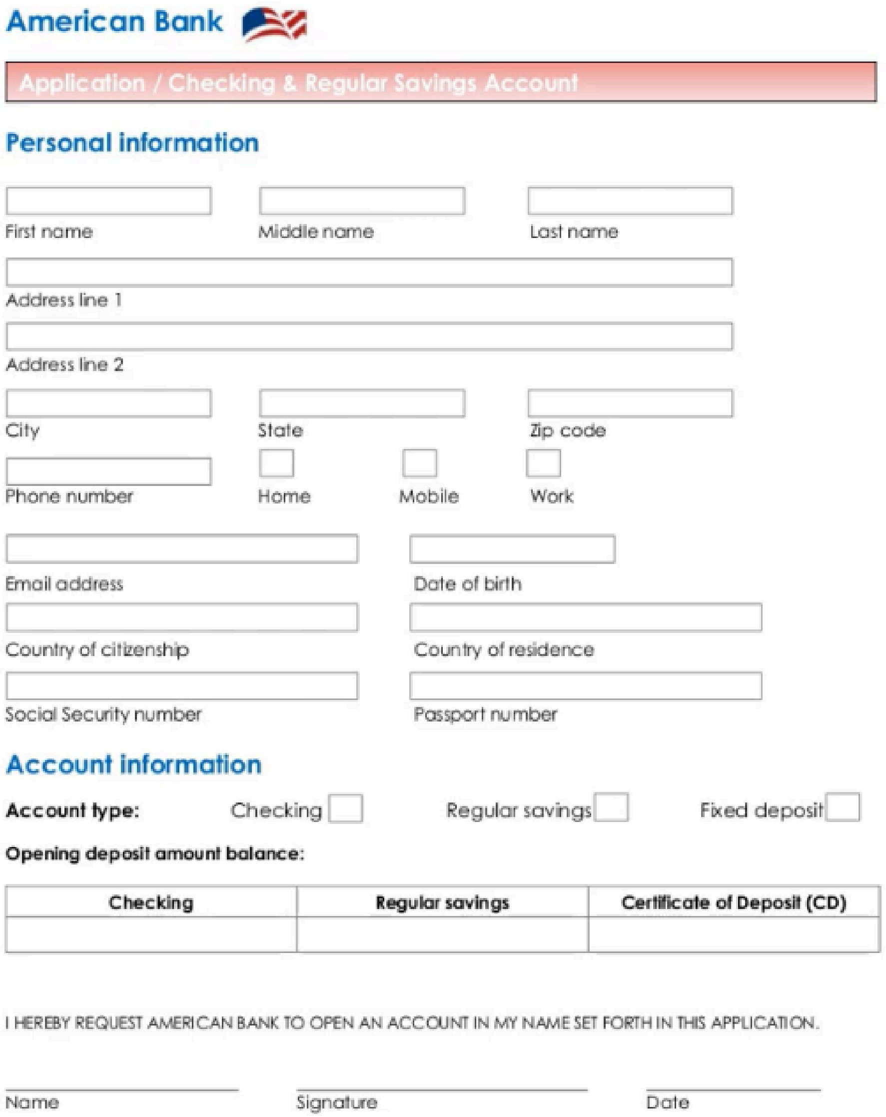

American Bank

Application / Checking & Regular Savings Account

Personal information

First name
Middle name
Last name
Address line 1
Address line 2
City
State
Zip code
Phone number
Home
Mobile
Work
Email address
Date of birth
Country of citizenship
Country of residence
Social Security number
Passport number

Account information

Account type: Checking Regular savings Fixed deposit

Opening deposit amount balance:

Checking	Regular savings	Certificate of Deposit (CD)

I HEREBY REQUEST AMERICAN BANK TO OPEN AN ACCOUNT IN MY NAME SET FORTH IN THIS APPLICATION.

Name
Signature
Date

Capital Rising, Investment Management & Market Provisioning are more related to NLP

- Banking Sector is the Most Affected
- Common People are Target audience
- Provide Automated Service for Target Audience

---

example of FinTech improving Service

Physical Bank to Banking Everywhere

- **Bank 1.0 –** Physical Bank
  - Over-the-counter Service
- **Bank 2.0 –** Online Bank
  - **Some** services can be done via the **Web**
- **Bank 3.0 –** Mobile Bank **(Where we are)**
  - Moneyless & Mobile Wallet
  - **Most** services can be done with **Smartphone**
- **Bank 4.0 – Banking Everywhere (Where we are going to be)**
  - **Technology-Enabled Innovation**
    - No Time Restriction & No Geographical Restriction
    - Blockchain
  - **AI-Enabled Innovation**
    - Precision Marketing (Personalized Recommendation System)
    - Real-Time Wealth Management Suggestion

embedded banking & emotional banking are the mechanisms that allow `Bank 4.0` to work

embedded banking

- key part of `Bank 4.0` b.c remove the boundary between financial and non-financial platforms
- financial services are integrated directly into non-financial platforms.
- use **Open APIs** (Open Banking), which make it easy to add financial functions to third-party platforms
- example : buying a game in a stream service

Emotional Banking : Proactive AI Assistance

- providing **"assistance for all about money"** exactly when and where a customer might need it

NLP in Finance

Two Traditional NLP research

Sentiment Analysis

Dictionary Construction

- build dictionaries that accurately reflect the "tone" of financial documents, as general-purpose sentiment dictionaries often fail in this context

mention 3 research papers related to Dictionary Construction

When Is a Liability Not a Liability? Textual Analysis, Dictionaries, and 10-Ks

- TIM LOUGHRAN, BILL MCDONALD
- **three-fourths (73.8%)** of the words flagged as negative by general dictionaries were actually neutral within business context.
- created the **Loughran-McDonald (LM) Financial Dictionary**, which remains an industry standard for calculating financial document sentiment today.

NTUSD-Fin: A Market Sentiment Dictionary for Financial Social Media Data

- market sentiment dictionary built from over 330,000 labeled posts on financial social media to evaluate investor attitudes.

Compare with SentiWordNet

Word	NTUSD-Fin	SentiWordNet	Word ID
buy	0.59	0.00	buy#1
sell	-0.98	0.00	sell#4
call	0.44	0.00	call#3
put	-0.49	0.00	put#1
overvalued	-3.42	0.25	overvalue#1
undervalued	1.21	-0.38	undervalue#1

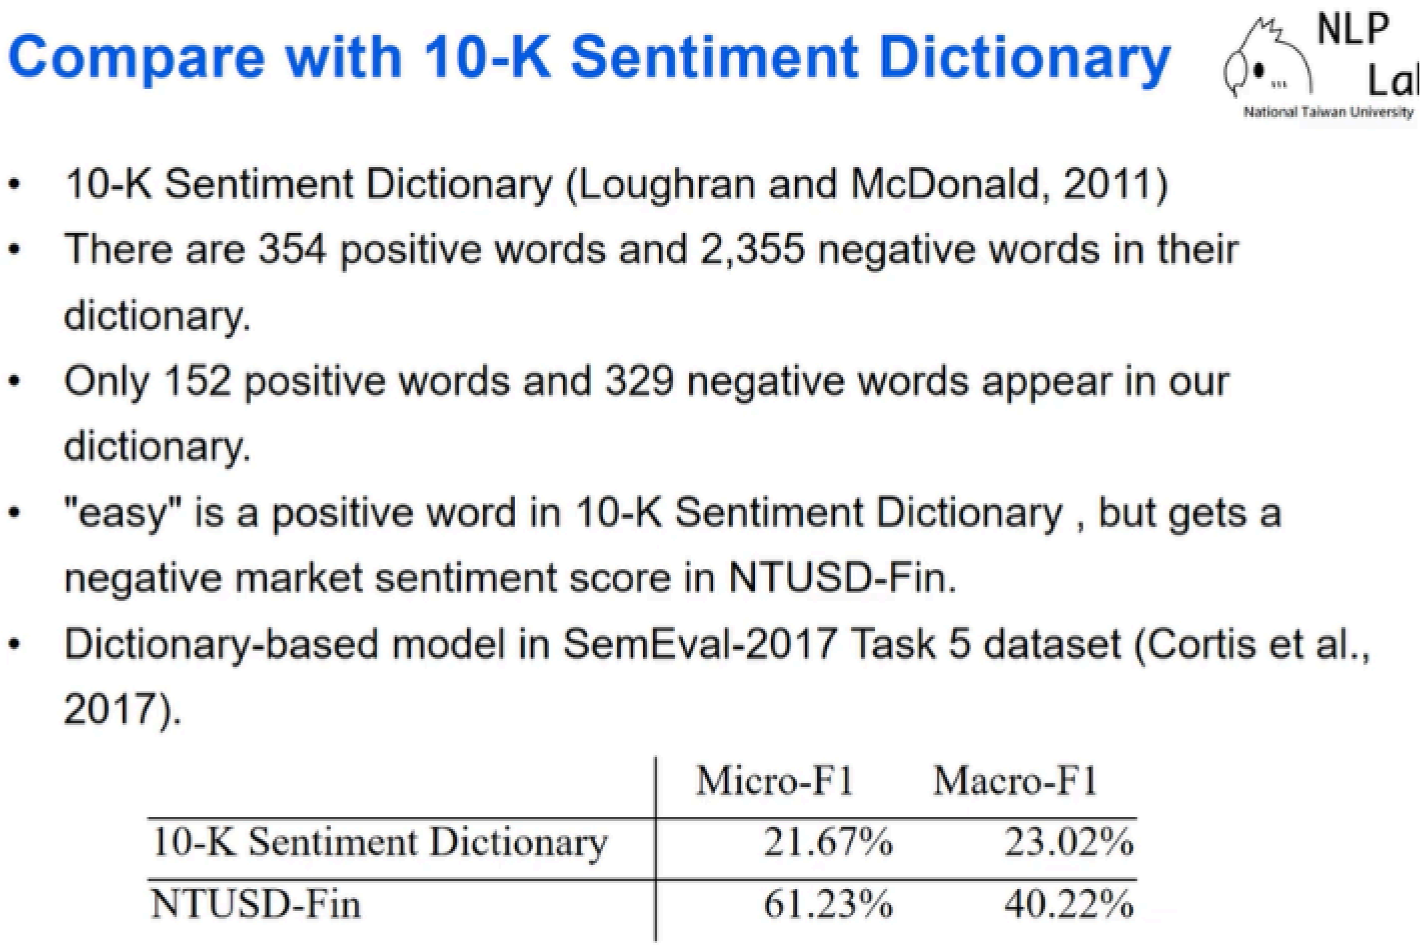

Compare with 10-K Sentiment Dictionary

NLP Lab
National Taiwan University

- 10-K Sentiment Dictionary (Loughran and McDonald, 2011)
- There are 354 positive words and 2,355 negative words in their dictionary.
- Only 152 positive words and 329 negative words appear in our dictionary.
- "easy" is a positive word in 10-K Sentiment Dictionary , but gets a negative market sentiment score in NTUSD-Fin.
- Dictionary-based model in SemEval-2017 Task 5 dataset (Cortis et al., 2017).

	Micro-F1	Macro-F1
10-K Sentiment Dictionary	21.67%	23.02%
NTUSD-Fin	61.23%	40.22%

Automatic Domain Adaptation Outperforms Manual Domain Adaptation for Predicting Financial Outcomes

- Core Finding : Automatically generated or adapted sentiment dictionaries perform better than manually created ones when forecasting financial market movements

Key Points

NLP Lab
National Taiwan University

- Automatically created sentiment dictionaries for predicting financial outcomes (Return & Volatility)
- 10-K reports 1994 to 2013
- Use the SVM to construct the dictionary (Threshold = 0.8)
- In empirical finance, a correct sentiment classification decision is not sufficient – the decision must also be interpretable and statistically sound. That is why we use ordinary least squares (OLS)
- The dependent financial variable is excess return or volatility.
- The value of the text variable for a category is the proportion of tokens from the category that occur in a 10-K.

$$y_i = \beta_1 x_{i1} + \beta_2 x_{i2} + \cdots + \beta_p x_{ip} + \varepsilon_i$$

Index Construction

use text mining to build indices that can **reflect or predict significant economic events** using keyword matching

Measuring Economic Policy Uncertainty

Findings

- Developing a new index of **economic policy uncertainty** (EPU) based on **newspaper** coverage **frequency**
  - Uncertainty or uncertain
  - Economic or economy
  - Congress, deficit, Federal Reserve, legislation, regulation or white house
- The index spikes near tight presidential elections, Gulf Wars I and II, the 9/11 attacks, the failure of Lehman Brothers, the 2011 debt-ceiling dispute and other major battles over fiscal policy.
- **Firm-level**: policy uncertainty raises stock price volatility and reduces investment and employment in policy-sensitive sectors like defense, healthcare, and infrastructure construction.
- **Macro-level**: policy uncertainty innovations foreshadow declines in investment, output, and employment in the United States and, in a panel VAR setting, for 12 major economies.

Economic Policy Uncertainty Index for Taiwan

- **Goal:** The index is designed to reflect or predict significant economic events by analyzing textual data.
- **Method:** It traditionally uses **keyword matching**, where researchers count specific keywords in newspapers to build a quantitative data point

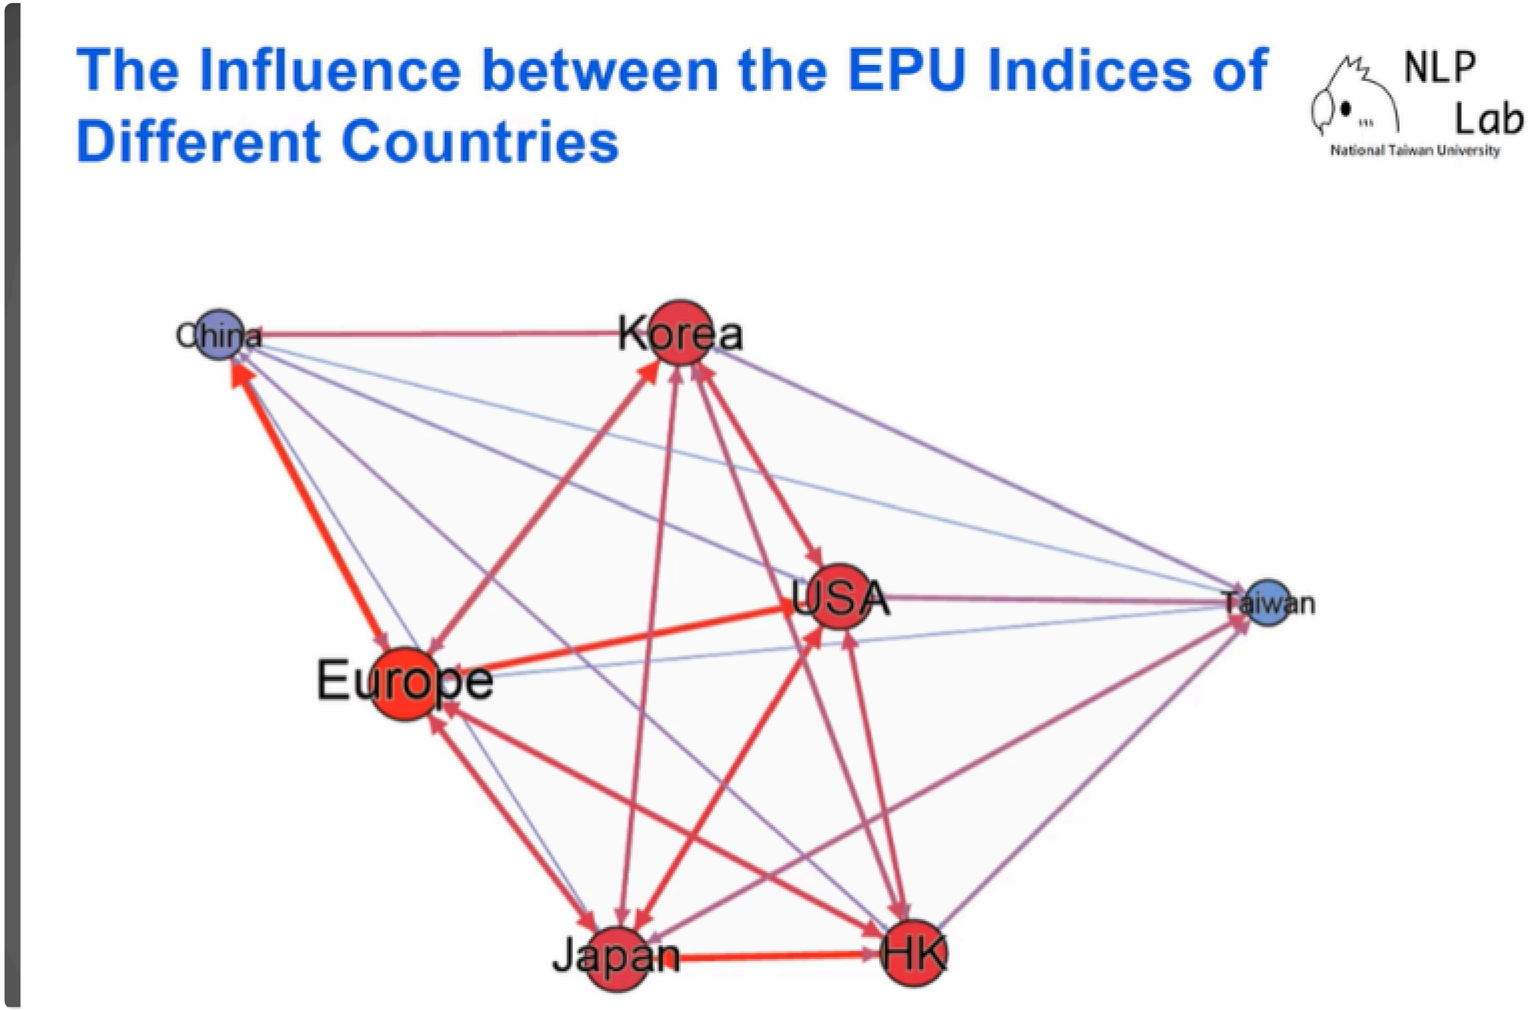

- **Findings** :
- The "Marginalisation" of Taiwan

Measure Country-level Socio-Economic Indicators with Streaming New: AN Empirical Study

- proposed automatically indicator construction method with the technique of event extraction & NER

Approach - Event-Centric Indicator Measure

NLP Lab
National Taiwan University

- Event Extraction
  - Extracts the predicate as an event trigger by the result of Semantic Role Labeling
  - Corpus-based event trigger extension & Check by human
- Location & Time
  - GeoNames & Timex2
  - Rule-based inference

Event triggers	Location	Time	Related*
the U.S. **recession** began in December 2007	U.S. (USA)	December 2007 (2007-12)	Y
SEPT. 1, 2011 The **bankruptcies** of three American solar power companies in the last month	American (USA)	Last month (2011-08)	Y
The man accused of fatally **shooting** his estranged wife inside a New Jersey church last Tuesday	New Jersey (USA)	Last Tuesday (2009-07)	N
...	...	...	

Indicator	Event triggers	Text
Unemployment rate	*unemploy, layoff, dismissal, lay off, economic crisis, bankrupt*	*Apple to* **lay off** *190 employees from self-driving car unit.*
CBOE Volatility Index (VIX) [6]	*sell share, liquidate, scandal, debt, loan* & (events for EPU)	*Humana executives* **sell** *nearly 74000 shares worth $22.7 million.*
Economic Policy Uncertainty (EPU)	*economic slowdown, institutional weaknesses, war, crisis, terrorist attack*	*These policies and other* **institutional weaknesses** *continue to undermine prospects for sustained economic development*

NLP in FinTech Applications

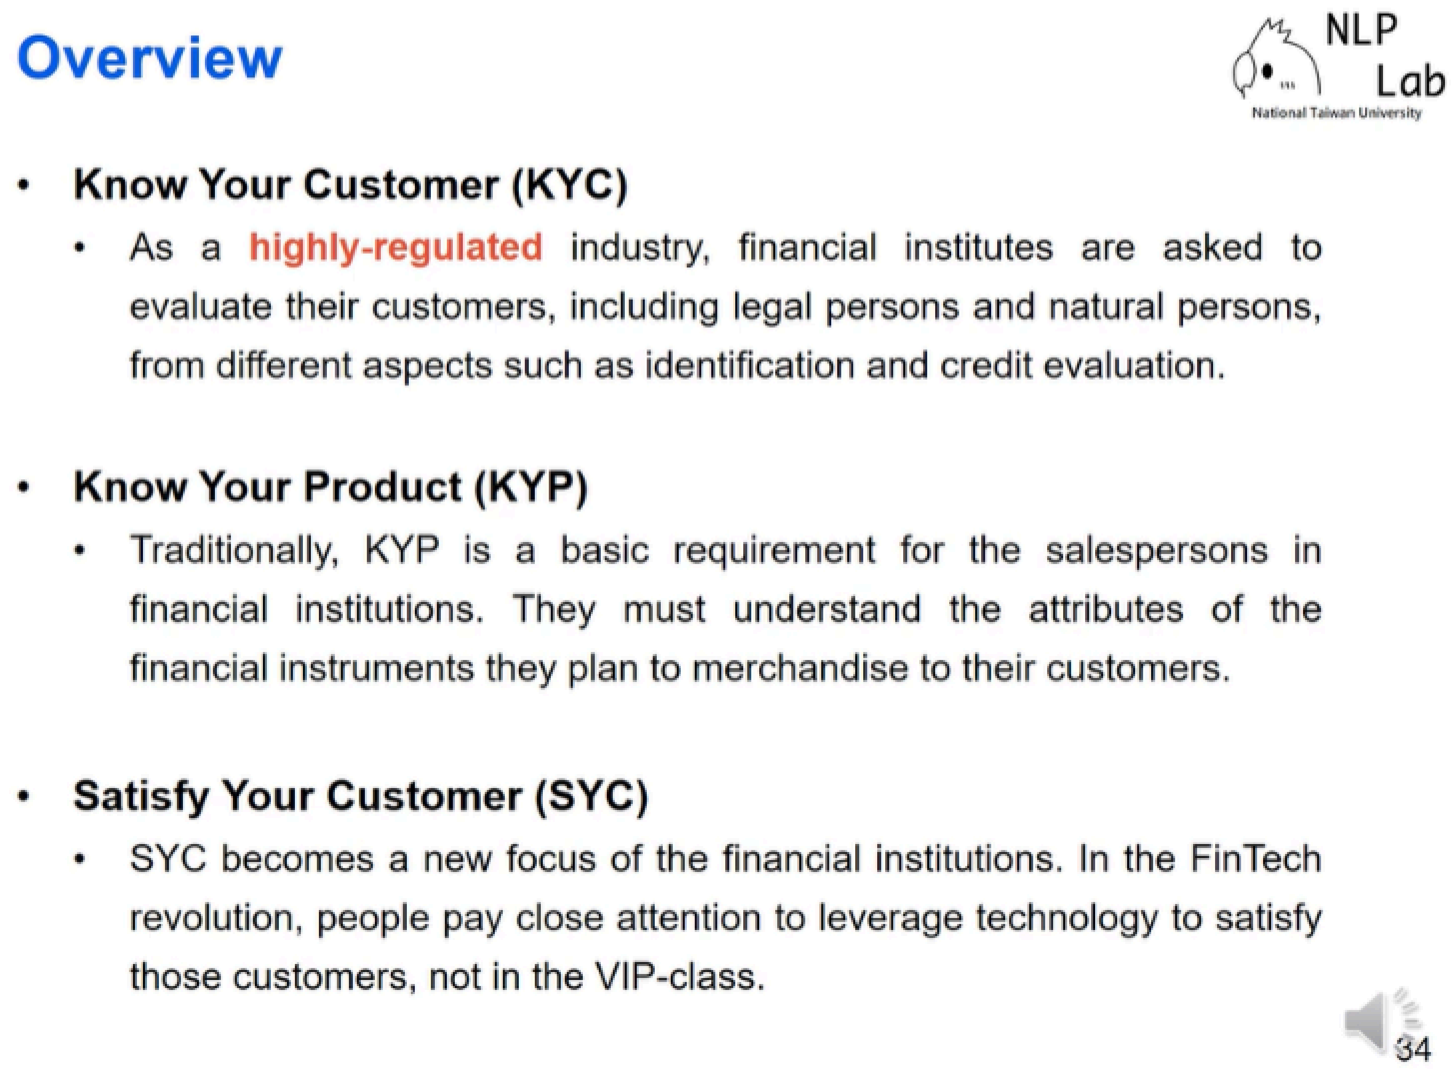

Know Your Customer (KYC)

a fundamental process necessitated by the highly regulated nature of the financial industry, requiring institutions to evaluate both legal persons (corporations) and natural persons (individuals) across various aspects such as identification and credit evaluation

KYC - Cooperate Customer

- **Past**
  - Event Extraction
  - Event Summarization
- **Future**

    - Event Forecasting
    - Identifying Predictive Causal Factors
    - Causality Detection
- **Relation**
    - Complex Relation Extraction
    - Network Embedding
- **Other Scenarios**
    - Exaggerated Numeral Information Detection
    - Legal Issues
    - Risk Prediction

Financial institutions use NLP to summarize a company's past events when first time facing.

Event Extraction

- converting past event unstructured documents into structured tables

paper :

- Doc2EDAG: An End-to-End Document-level Framework for Chinese Financial Event Extraction ->
    - End-to-end document-level financial event extraction.

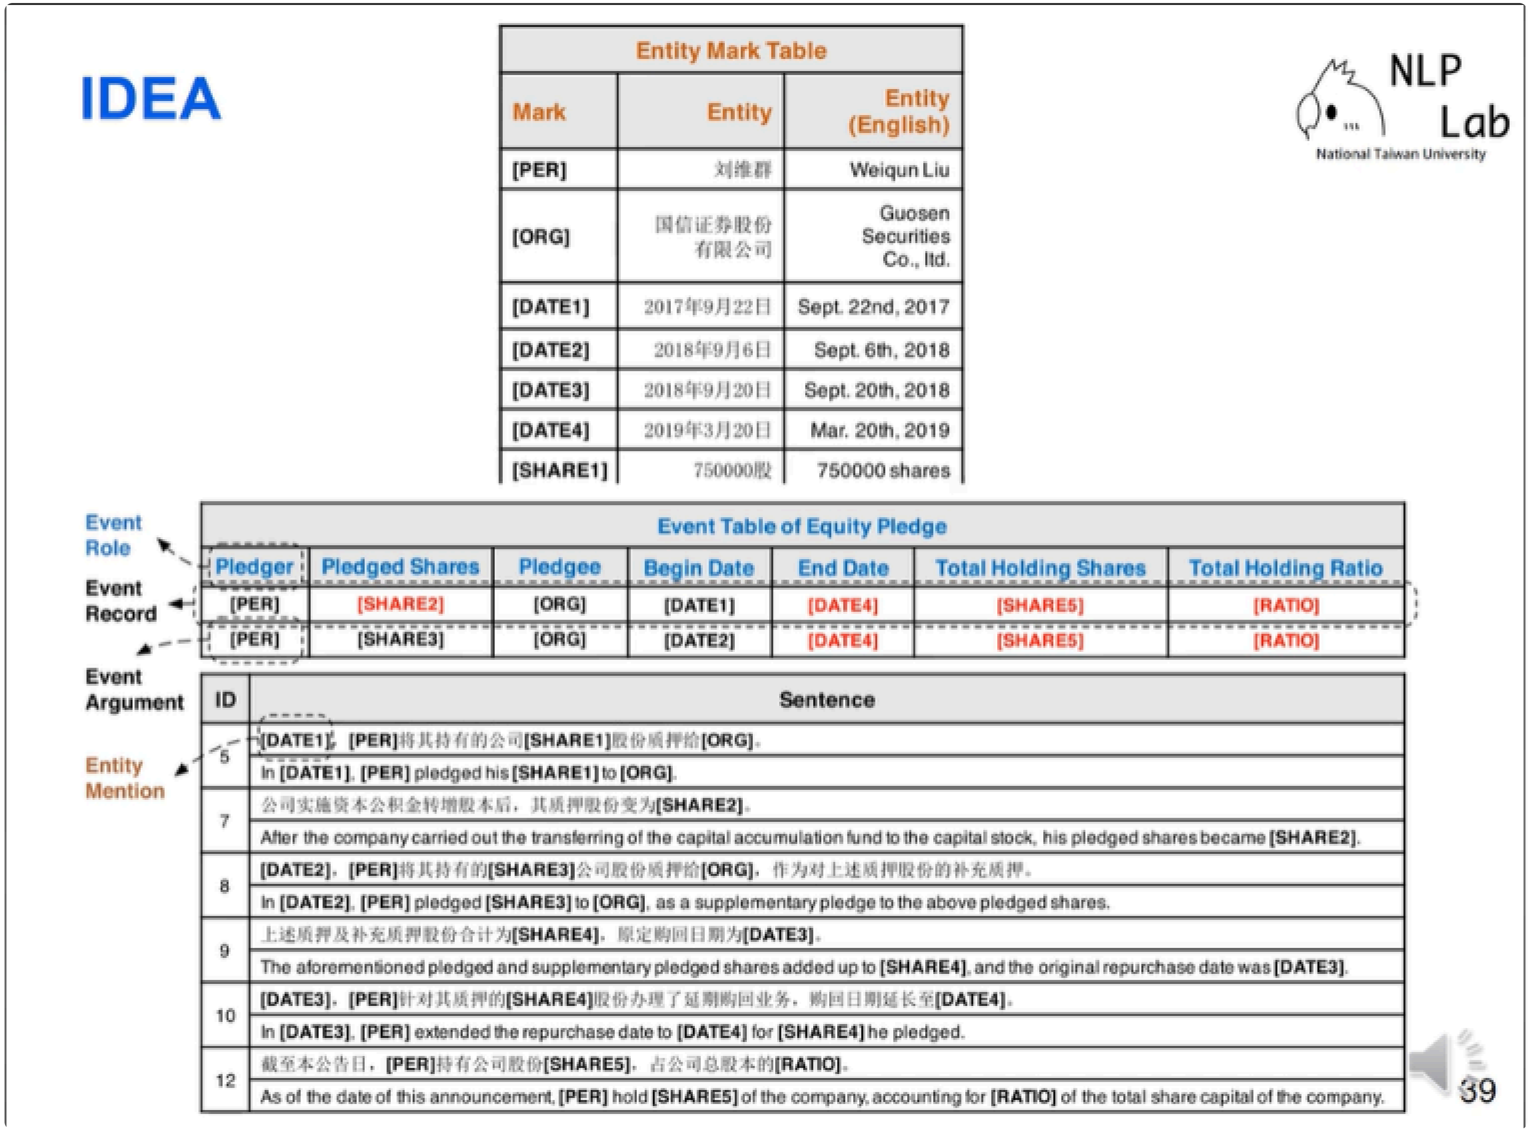

Event Summarization

paper :

- A Unified Model for Financial Event Classification, Detection and Summarization

Results

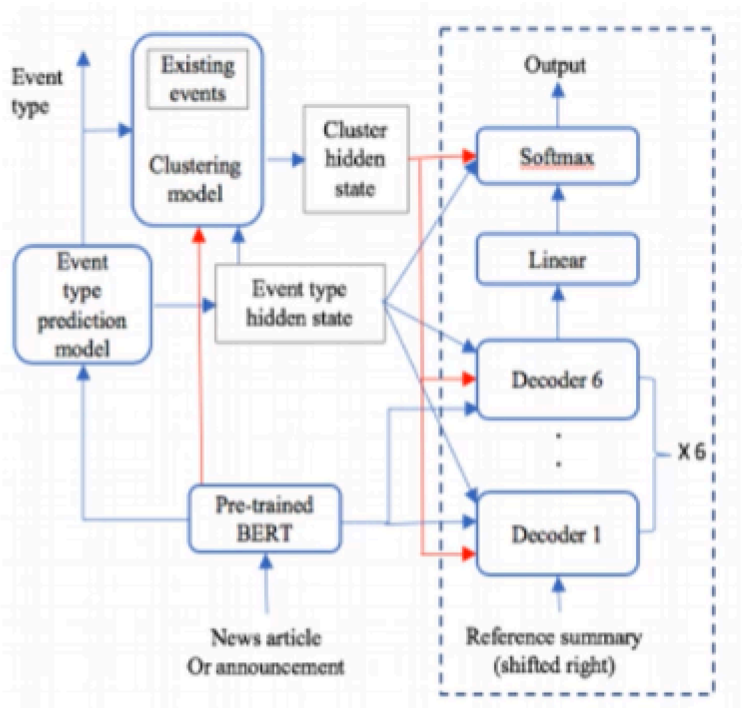

Method	F1
HAN	92.5
ULMFit	93.3
BERT	94.8
Our approach	95.5

Table 2. Event type classification result

Method	NMI	B-Cubed
UMass	0.665	0.310
LSH	0.613	0.263
SemEntity	0.708	0.391
BERT-sim	0.717	0.397
Our approach	0.742	0.414

Table 3. Event clustering result

Method	R-1	R-2	R-L
Bi-LSTM-attention	52.8	34.0	49.2
DCA	54.1	35.5	50.9
Our approach	56.5	37.7	53.3

Table 4. Event summarization result

Future Event Forecasting

papers :

- Forecasting Firm Material Events from 8-K Reports
  - Uses SEC filings (8-K reports) to predict future important corporate events.

Causality Detection

focuses on establishing directional, cause-and-effect relationships between variables rather than simple correlations (, not just)

papers

- NTUNLPL at FinCausal 2020, Task 2: Improving Causality Detection Using Viterbi Decoder
  - -> Build NLP models specifically for detecting causal relationships inside financial text.

Identifying Predictive Causal Factor

both **causality detection** & **Identifying Predictive Causal Factor** deal with understanding relationship between variables, but focus is different

---

Identifying Predictive Causal Factor
isolates the specific variables that directly drive and reliably forecast an outcome,

papers

- Identifying Predictive Causal Factors from News Streams
  - -> find cause from news stream

Complex Relation Extraction

identifies implicit and explicit relationships between companies and industries through transaction records, news, and contact histories

papers :

- Extracting Complex Relations from Banking Documents
  - Extracting Complex Relations from Banking Documents
- Financial Text Data Analytics Framework for Business Confidence Indices and Inter-Industry Relations

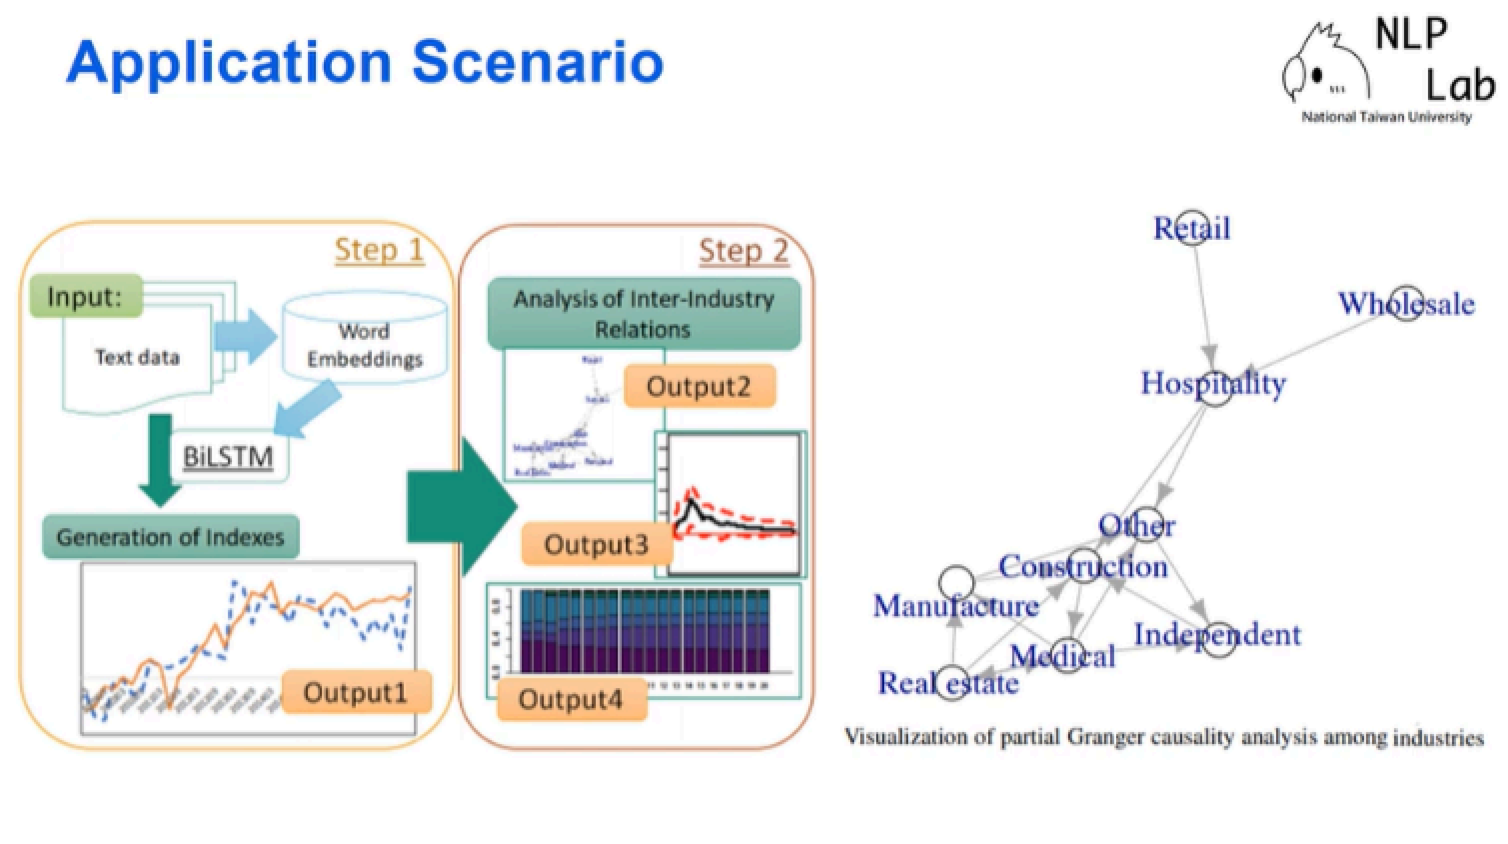

Network Embedding

representing companies as vector

papers

- News2vec: News Network Embedding with Subnode information
  - Build graph embeddings of a company from financial news.

Exaggerated Numerical Information Detection

paper

- Numeracy-600K: Learning Numeracy for Detecting Exaggerated Information in Market Comments
  - Created a 600K dataset for numerical reasoning in financial text

Legal Issues

Identifies active litigation

papers

- Step-Wise Refinement Classification Approach for Enterprise Legal Litigation
  - A step-wise refinement classification approach for enterprise legal litigation replacing Black-Box (NN)
- e-Commerce and Sentiment Analysis: Predicting Outcomes of Class Action Lawsuits
  - demonstrating that the **textual sentiment in corporate SEC disclosures (MD&A and Market Risks sections) can accurately predict the outcomes of securities class action lawsuits** against e-commerce companies

Risk Prediction

- evaluate the operation of a company for a customer oriented product

papers

- Learning to Learn Sales Prediction with Social Media Sentiment
  - Transferring knowledge from source products significantly reduced sales prediction error for new target products compared to traditional time-series and baseline models.
- Interpretable Operational Risk Classification with Semi-Supervised Variational Autoencoder
  - presents **SemiORC**, a semi-supervised text classification framework that addresses two major hurdles in automated Operational Risk Management (ORM): **severe data imbalance (sparse labels)** and a **lack of model interpretability**.

KYC - Personal Customers

Unlike corporate KYC (understanding companies), **Personal KYC** focuses on understanding an individual customer.

traditional KYC relies one

- application forms
- credit reports

NLP for personal KYC

Writing Style Analysis

paper

- You Write Like You Eat: Stylistic variation as a predictor of social stratification
  - analyze a customer's **writing style** to predict their **social stratification** (social class), which helps in risk and credit assessment

personal knowledge graph

Instead of treating every customer as isolated information, build a graph connecting everything known about them.

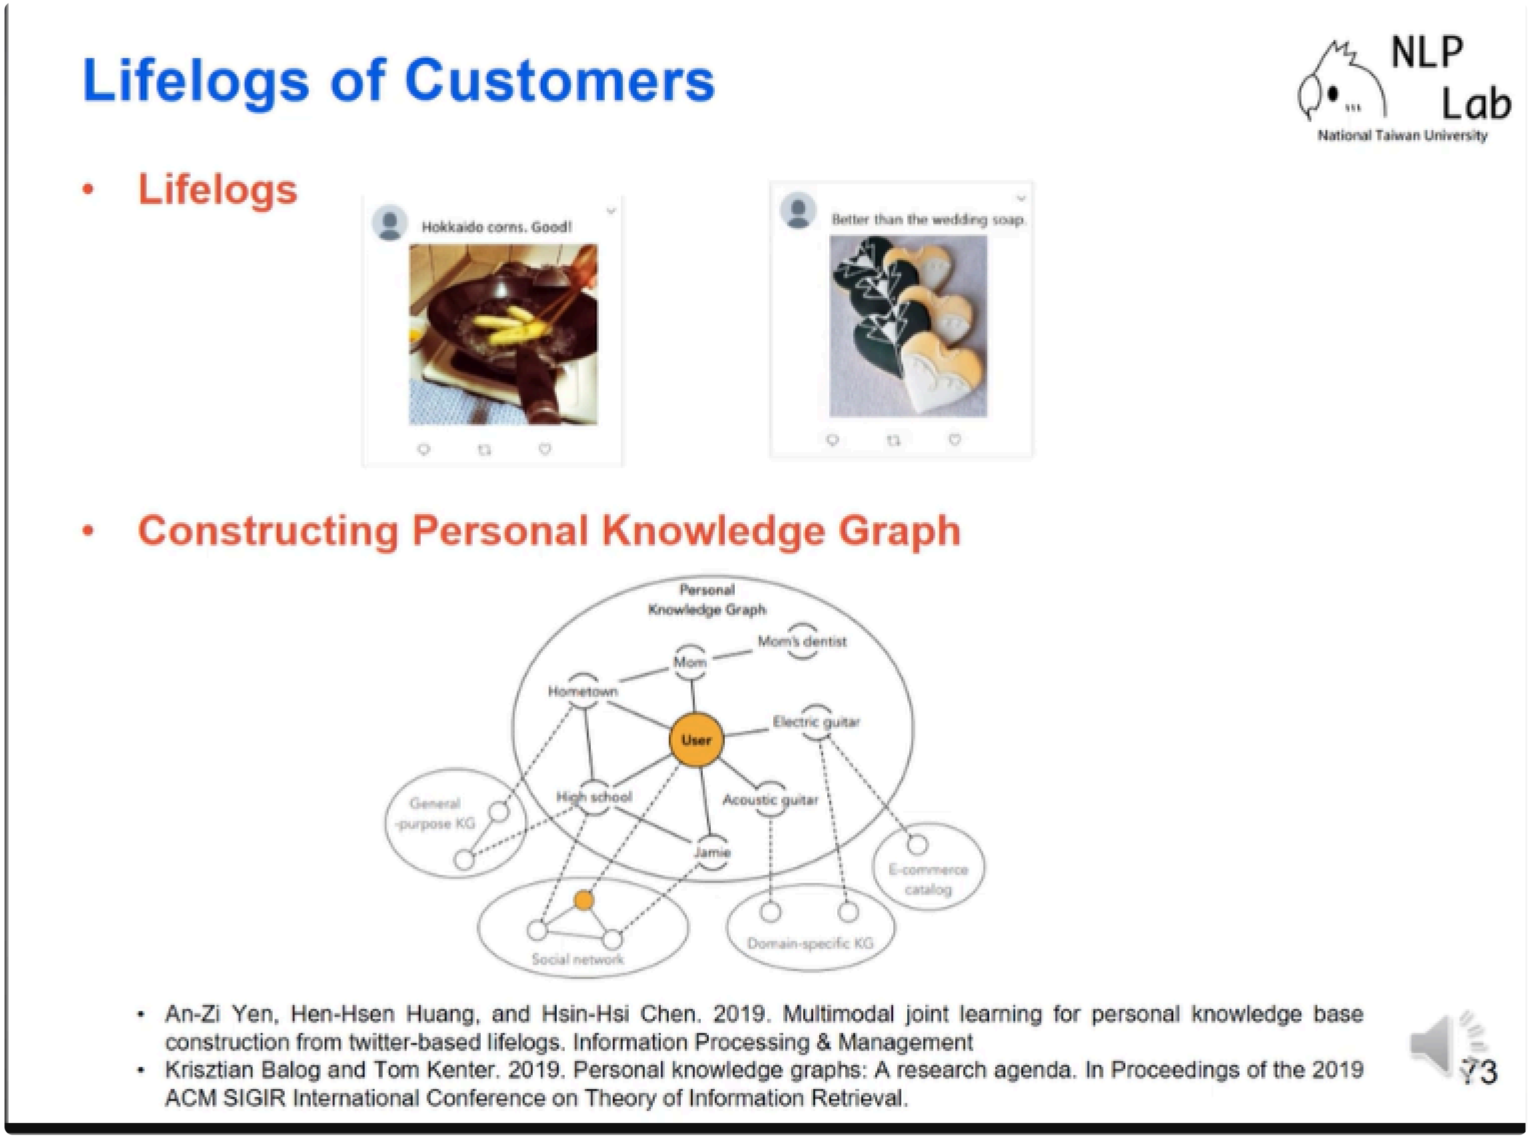

Early Detection

- **Credit Problem**
  - Start attacking the problem earlier
- **Diseases**

- The insurance companies can encourage and support their customers to get treatment early
- Benefit to both insurance institute and their customers

Know Your Product (KYP)

Traditionally, KYP meant that financial advisors had to understand the basic characteristics of products (stocks, bonds, insurance, mutual funds) before recommending them.

In modern FinTech, KYP becomes much broader:

- continuously updating product information
- estimating future prospects
- evaluating risk
- explaining recommendations with AI

in here, product can be stock

Task	Information Source	Goal	Main Contribution
Earnings Call Analysis	Transcript + Voice	Predict stock risk	Introduced multimodal (text + audio) financial prediction
Earnings Call Q&A	Analyst–Executive dialogue	Predict financial risk	Showed Q&A is more informative than prepared speeches
Analyst Behavior	Analysts' questions	Predict changes in analyst forecasts	Used semantic + pragmatic features to model analyst decision making
Numeral Understanding	Financial tweets	Build market indicators	Treated numerals as structured financial information instead of plain text

Earnings Conference Call Analysis

An earnings conference call

- a live audio meeting where a public company's top leaders discuss their financial results, share future goals, and answer questions from market experts
- contain text, audio data

papers :

- What You Say and How You Say It Matters: Predicting Financial Risk Using Verbal and Vocal Cues
    - used to analyze both the **transcripts (textual data)** and **vocal cues (audio data)** from these calls to predict stock price volatility
    - can notice even if Text looks positive But voice sounds hesitant.
    - Predicts **future stock volatility (financial risk)**

Risk Prediction from Earnings Call Q&A

the Question & Answer section is often more informative than prepared speeches.

papers :

- Financial Risk Prediction with Multi-Round Q&A Attention Network
    -
        - Introduced attention over multiple Q&A turns.
    - Showed Q&A contains stronger predictive signals than prepared speeches.
    - Improved financial risk prediction.

Analyst Behavior Analysis

predict changes in financial forecasts by analyzing how professional analysts interact with company leaders during Q&A sessions

papers :

- Modeling Financial Analysts' Decision Making via the Pragmatics and Semantics of Earnings Calls
    - analyze the **programmatic and semantic features** of the questions analysts ask to understand their underlying decision-making process

    - to predict **changes in analyst forecasts**, such as adjustments to a company's "price target"
    - **finding** : specific analysis ask into more specific direction

Understanding Numbers in Social Media

utilizes **social media data** (such as Twitter) to extract numerals and construct indicators. Research shows that investors on social media ("the crowd") tend to be more **progressive** with their price targets compared to professional analysts, though following them can yield similar results

papers :

- Fine-grained Numeral Understanding in Financial Tweet

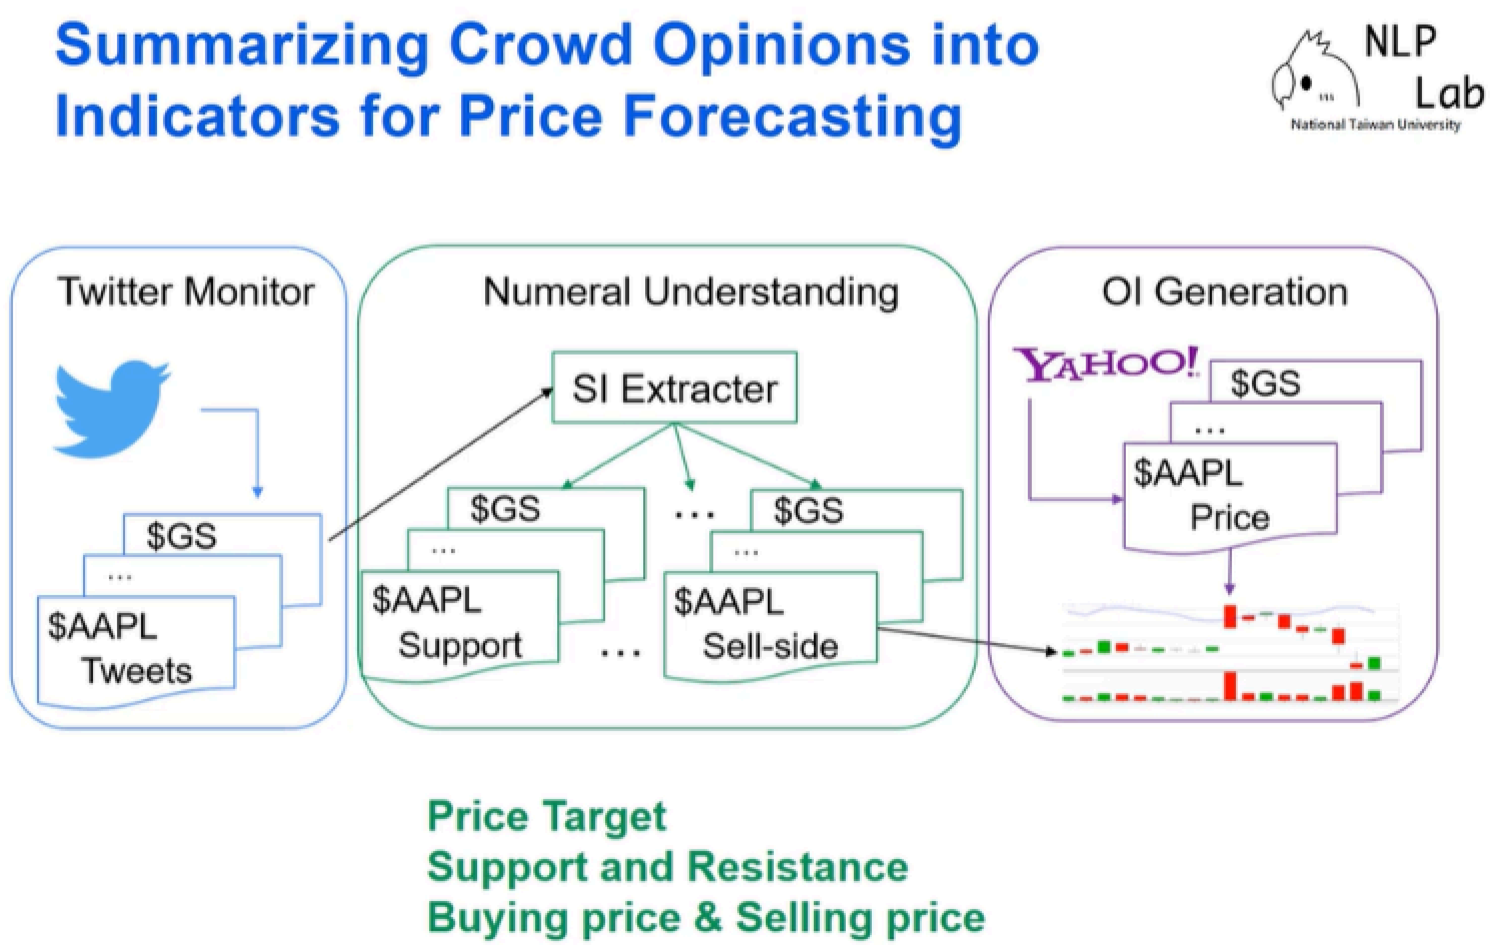

- extract and classified number into categories
- release **The FinNum Benchmark Dataset:**

Satisfy Your Customer (SYC)

the most important issue in the FinTech

Topic	Main NLP Task	Main Contribution
Wealth management	Recommendation	Combine news and prices to recommend portfolios

Topic	Main NLP Task	Main Contribution
Story ending prediction	Narrative reasoning	Shows narrative prediction methods can inspire financial forecasting and advice
NClaim	Claim detection	Detect investment recommendations in analyst reports
Rational classification	Argument mining	Evaluate the quality of evidence supporting investment advice
Analyst report quality	Text classification	Predict whether analyst reports are reliable
Explainable recommendation	Explainable AI	Recommend products together with understandable reasons
Financial content recommendation	Recommendation systems	Personalize financial news and information
Insurance recommendation	Knowledge graph recommendation	Recommend insurance products for cold-start users

Wealth Management Suggestion

Help customers make investment decisions automatically.

papers

- Stock Embeddings Acquired from News Articles and Price History, and an Application to Portfolio Optimization
    - The model combines stock prices & financial news to recommend investment portfolios

inside a neural network

to recommend investment portfolios.

Story Ending Prediction → Financial Advice

Story Ending Prediction is a task that needs to select an appropriate ending for the given story, which requires the machine to understand the story and sometimes needs common-sense knowledge.

Professional analysts perform almost the same reasoning.

They

- summarize facts
- understand events
- predict future outcomes

papers

- Discriminative Sentence Modeling for Story Ending Prediction

NClaim Dataset

papers

- NumClaim: Investor's Fine-grained Claim Detection

Motivation

NLP Lab
National Taiwan University

- Over **58.47%** of sentences in analysis report contain at least one numeral

- Investors always make a claim **with an estimation**
  - **(X) We estimate that the sales may growth**
  - **(O) We estimate that the sales growth rate may exceed 40%**

- The importance of fine-grained claims and the numerals.
  - We estimate that the sales growth rate may exceed 5%
  - We estimate that the sales growth rate may exceed 40%

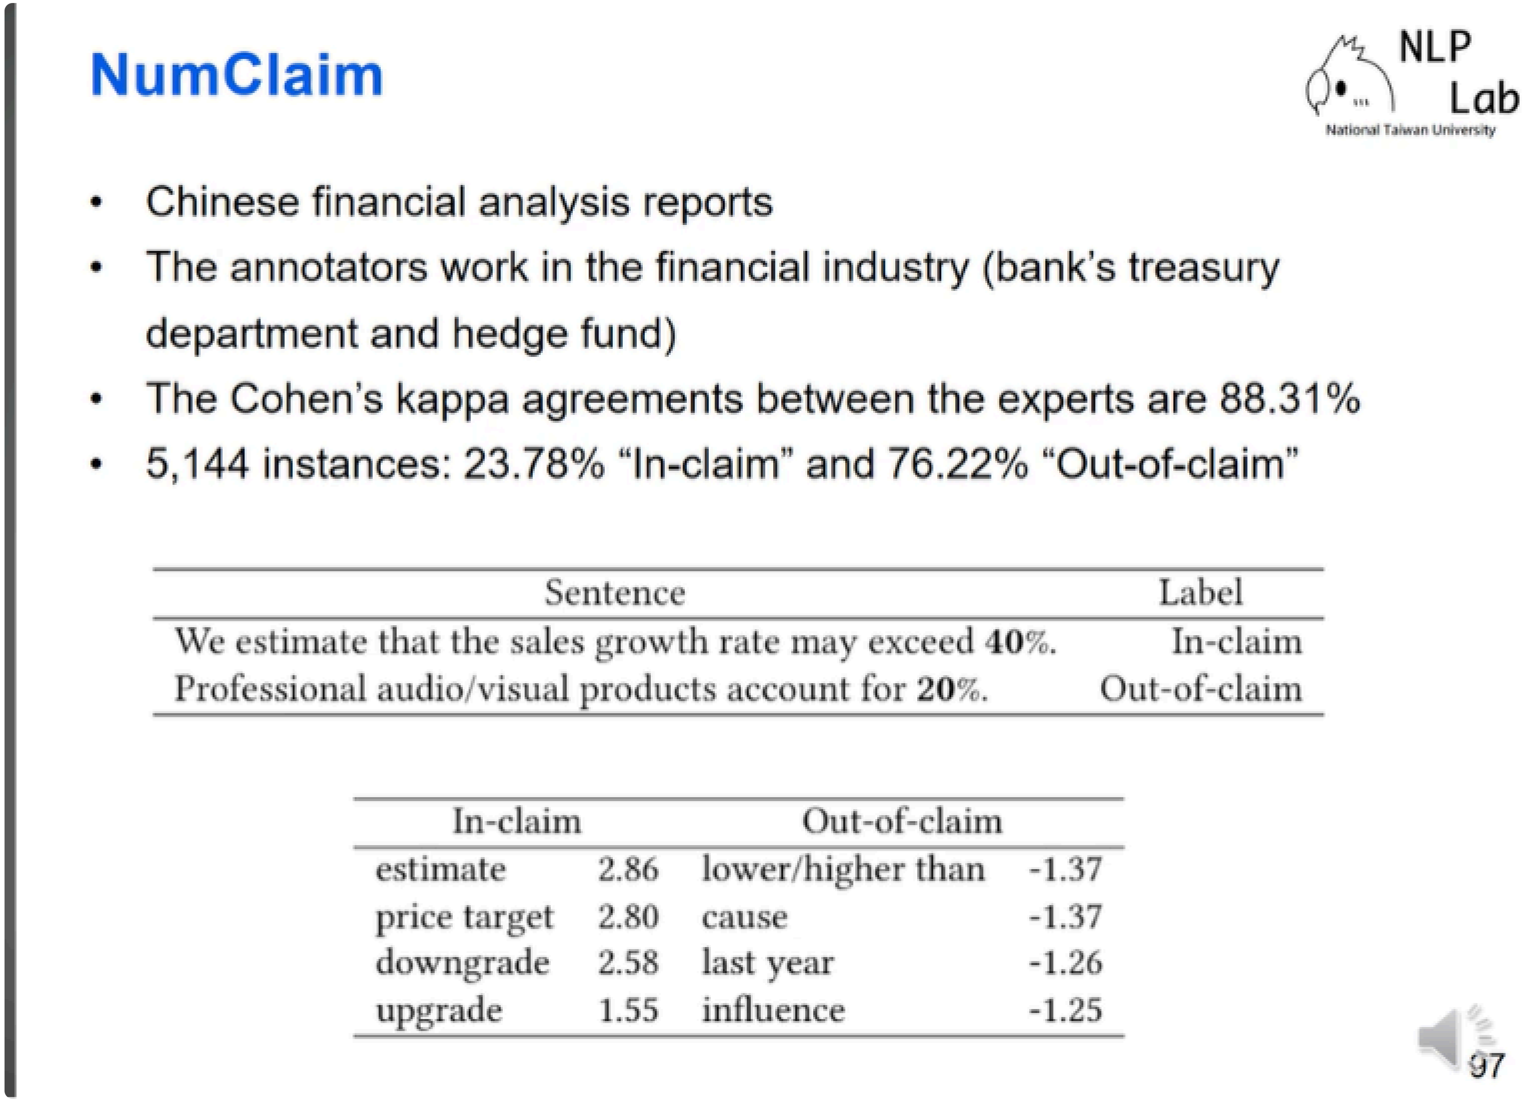

- Designed for argument mining in the financial domain to identify and extract key claims made by investors.
- Published with expert annotations, distinguishing between in-claim and out-of-claim numerical mentions (frequently used in tasks like FinNum)

Investment Rationale Classification

After detecting claims, the next step is evaluating the supporting evidence.

papers :

- Rationale Classification for Educational Trading Platforms

IDEA

NLP Lab
National Taiwan University

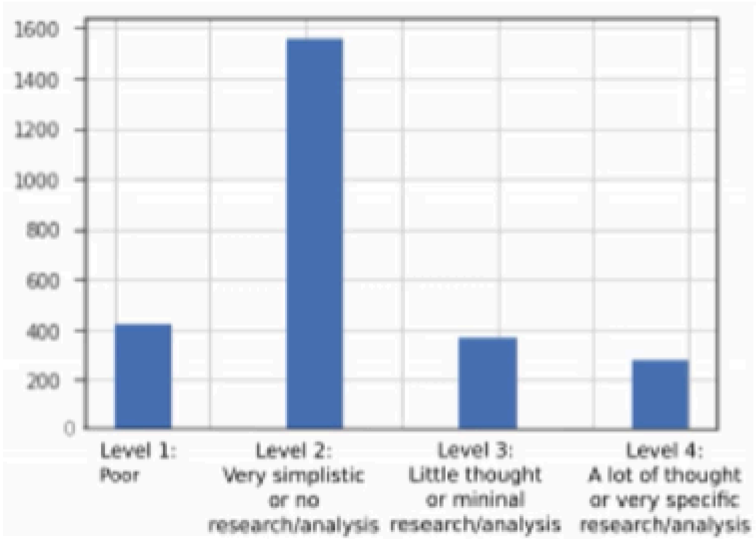

Figure 1: Distribution of rationales by levels of thoughtfulness

1. A rationale contains little or no thought.
2. A rationale contains research or analysis that is simplistic or too general.
3. A rationale contains specific research or analysis.
4. A rationale contains an extensive amount of research or analysis.

Chen, Chung-Chi, Hen-Hsen Huang, and Hsin-Hsi Chen. "Fine-grained Financial Opinion Mining: A Survey and Research Agenda." *arXiv* (2020): arXiv-2005.

99

-

    - Automates the classification of whether a student's rationale is thoughtful or not using convolutional neural networks (CNNs) and support vector machines.

Accurate vs Inaccurate Analyst Reports

papers

Measuring Forecasting Skill from Text

- *forecasting skill can be accurately measured directly from the language used in predictions*, proving that high-accuracy forecasters exhibit distinct linguistic patterns (such as calibrated uncertainty, higher readability, and neutral tone) across geopolitical and financial domains.

Explainable Recommendation

Traditional recommender : recommend stock A -> no explanation

Explainable recommender : recommended Stock A -> lower risk / similar return / fits your profile

papers

Try This Instead: Personalized and Interpretable Substitute Recommendation

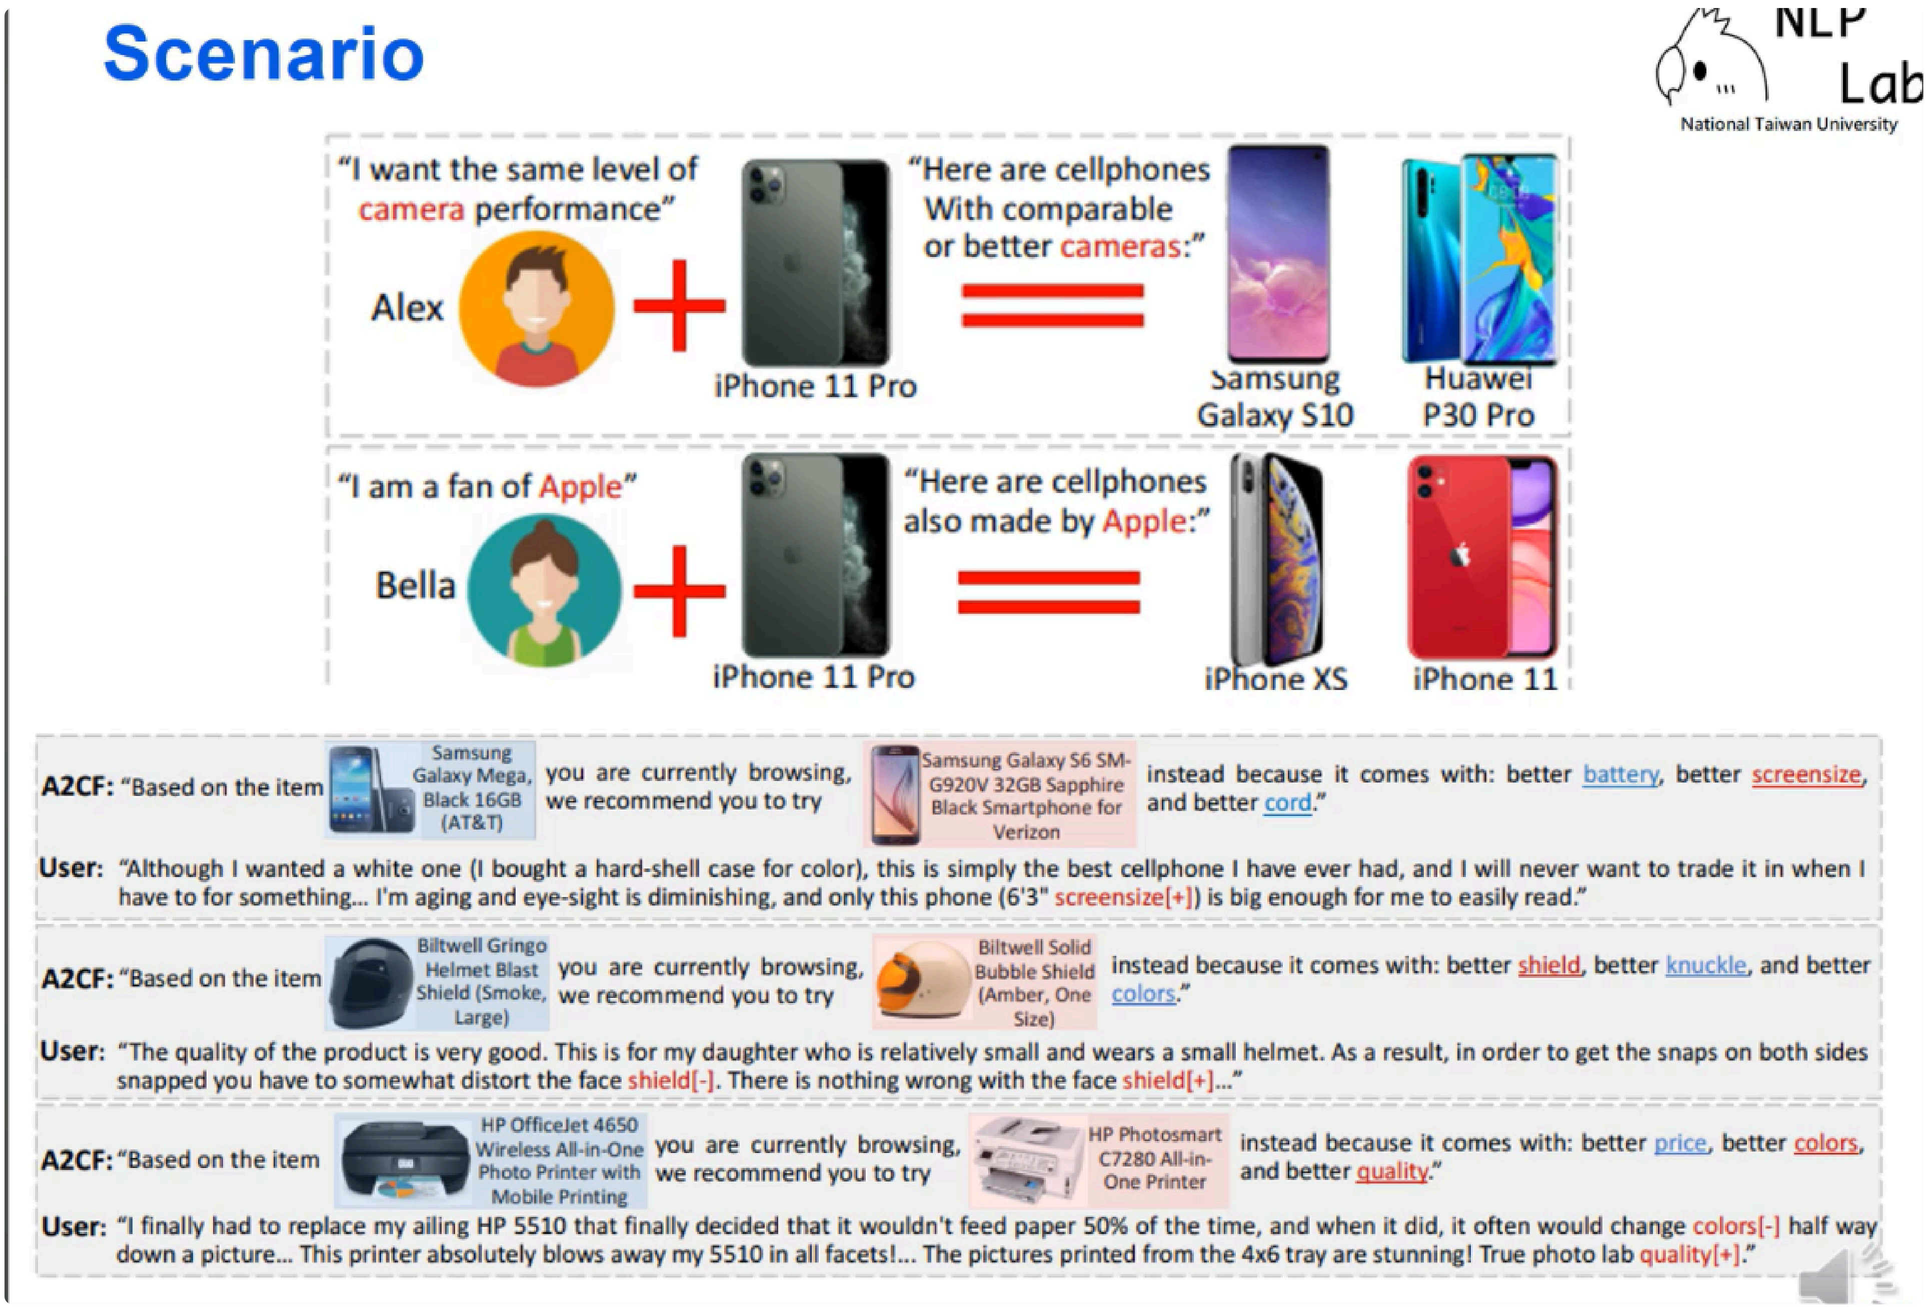

- Generate recommendations with human-understandable explanations.

Personalized Financial News Recommendation

Every investor has different interests.

papers :

Next Cashtag Prediction on Social Trading Platforms with Auxiliary Tasks

- introduces **personalized next-cashtag prediction** on social trading platforms by combining textual BERT representations with stock price trends, using an **Attentive Capsule Network** trained jointly with **market-driven auxiliary tasks** (predicting mention growth and stock profitability).
- By combining **what people write**, **how stock prices move**, and **what's trending**, the AI gets really good at guessing the exact `$CASHTAG` you'll want to talk about next!

ATBRG: Adaptive Target-Behavior Relational Graph Network for Effective Recommendation

- constructs a **dynamic, pair-specific sub-graph** bridging user behaviors and the target item instead of representations over the entire Knowledge Graph (which introduces noise and high computational cost).

Personalized Insurance Recommendation for new comers

Insurance differs from online shopping.

Problems

- rarely purchased
- complicated products
- little historical data

Paper

DCDIR: A Deep Cross-Domain Recommendation System for Cold Start Users in Insurance Domain

- Use
  - insurance knowledge graph
  - user behavior from another domain
    to recommend appropriate insurance products.

Further Exploration

Crowdfunding

Crowdfunding is raising small amounts of money from many people online. Success depends heavily on *persuasiveness*

papers :

Let's Make Your Request More Persuasive: Modeling Persuasive Strategies via Semi-Supervised Neural Nets on Crowdfunding Platforms

Details

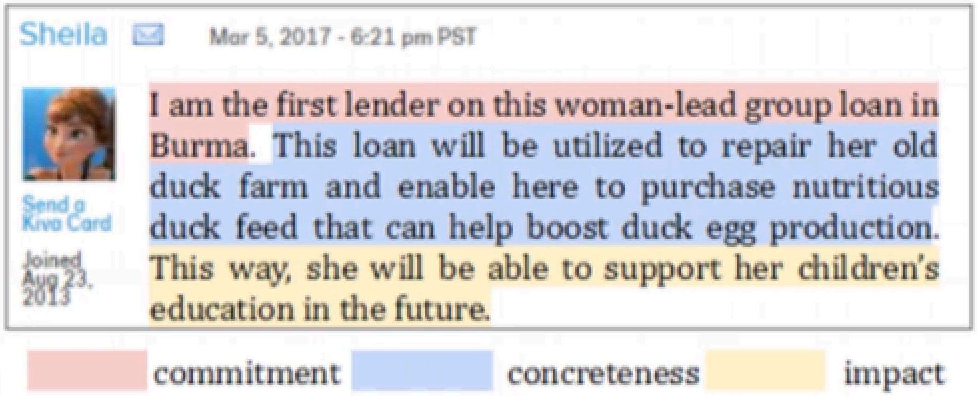

Strategy	Top Ranked Keywords
Commitment	joined, lenders, loaning, lend, loan just, join, loaned, made, lent
Concreteness	women, married, old, heads, year-old money, sells, years, business, number
Emotion	hard, thank, better, grief, great maybe, help, please, thanks, happy
Identity	promotion, shall, captain, form, number spirits, lenders, member, team
Impact	improve, new, better, products, money to, use, business, more, order
Scarcity	minutes, there's, now, soon, go expire, hours, days, number, left

- Builds a neural network to measure persuasiveness & identify persuasion strategies
- Using **word- and sentence-level attention**, the model not only predicts overall request success but also pinpoints which specific words and sentence tactics contribute most to persuasive outcomes across different crowdfunding platforms

The supervision signal is

> Number of lenders who funded the request.

FinBERT

Pretrained BERT on

- financial documents
- general text
  to create a domain-specific language model

papers :

- FinBERT: A Pre-trained Financial Language Representation Model for Financial Text Mining

Compliance Checking

Banks must follow many regulations. Manual compliance checking is

- slow
- expensive
- inconsistent.

Paper

Deep Semantic Compliance Advisor for Unstructured Document Compliance Checking

- Uses NLP to automatically compare regulations & company documents to detect compliance issues.
- use *Clause-Level Matching* to find the most relevant standard reference clauses for any given contract clause.
- use *Statement-Level NLI* to determines whether a contract statement semantically *entails* (complies with), *contradicts* (violates), or is *neutral* toward a reference standard statement
- Use **Graph Neural Network (GNN) syntactic sentence encoder** that builds dependency parse graphs of legal statements. This allows the model to capture deep syntactic relationships and long-distance dependencies much better than standard sequential encoders

Anti-Money Laundering (AML)

traditionally employees manually inspect customer records, transaction histories and news.

Paper

Pluto: A Deep Learning Based Watchdog for Anti Money Laundering

- a distributed NLP batch system that uses **paragraph embeddings and news clustering** to automate adverse media screening, reducing manual AML compliance workload by **67%**

Human Resource Automation

traditionally employees manually inspect customer records, transaction histories and news.

Paper

Learning to Ask Screening Questions for Job Postings

- propose a model to suggest screening question based on job posting

Scenario

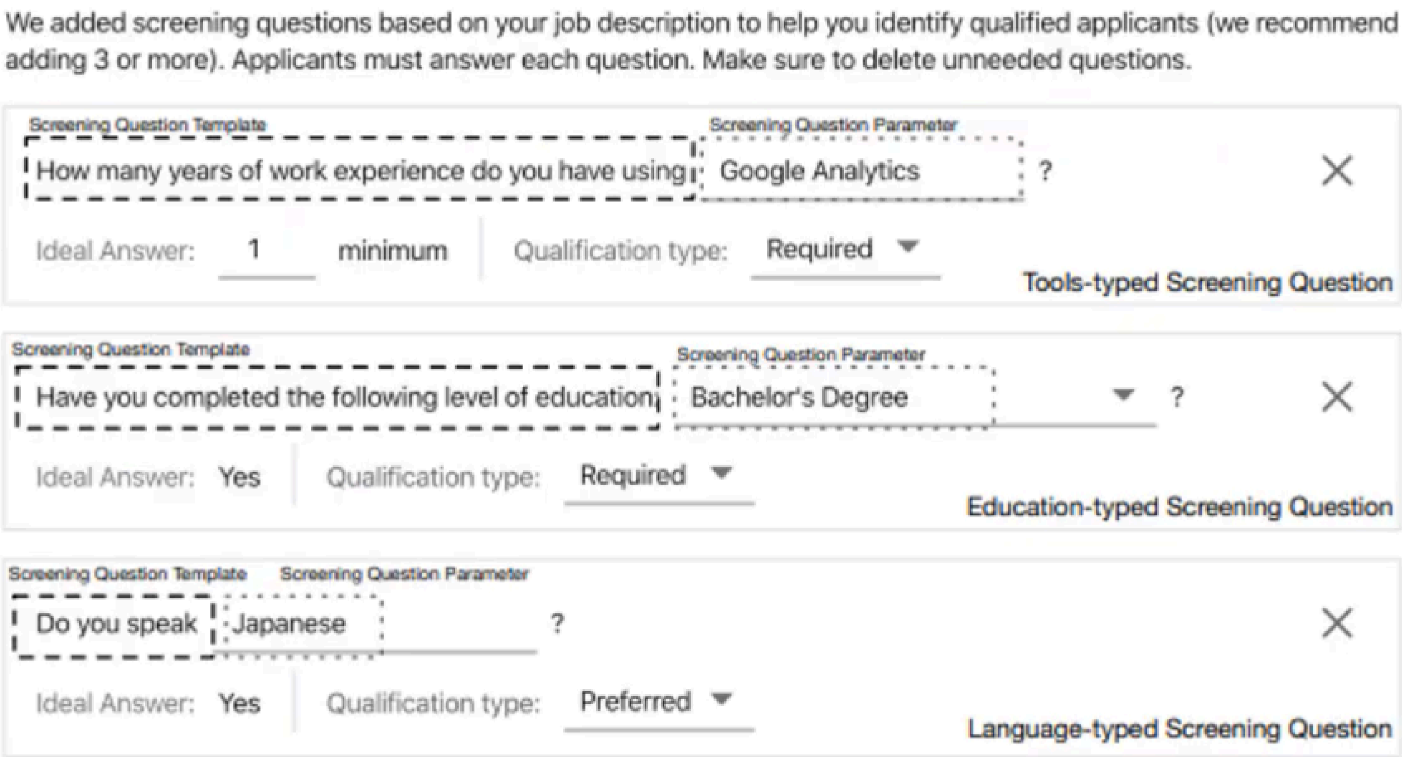

We added screening questions based on your job description to help you identify qualified applicants (we recommend adding 3 or more). Applicants must answer each question. Make sure to delete unneeded questions.
Screening Question Template
Screening Question Parameter
How many years of work experience do you have using
Google Analytics
?
Ideal Answer: 1 minimum
Qualification type: Required
Tools-typed Screening Question
Screening Question Template
Screening Question Parameter
Have you completed the following level of education
Bachelor's Degree
?
Ideal Answer: Yes
Qualification type: Required
Education-typed Screening Question
Screening Question Template
Screening Question Parameter
Do you speak
Japanese
?
Ideal Answer: Yes
Qualification type: Preferred
Language-typed Screening Question

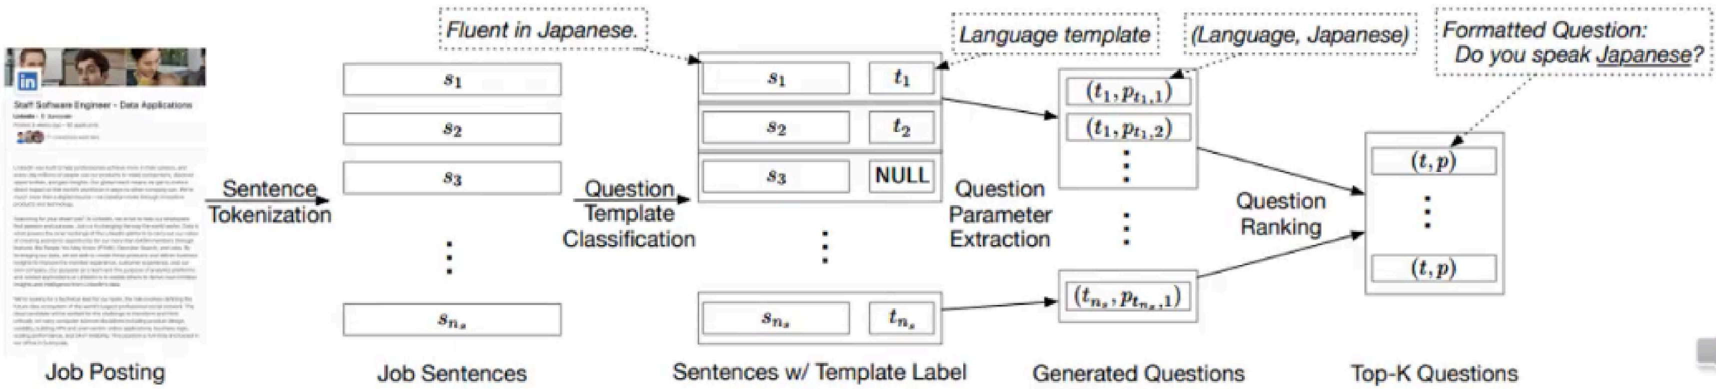

Fluent in Japanese.
Language template
(Language, Japanese)
Formatted Question: Do you speak Japanese?
Sentence Tokenization
Question Template Classification
Question Parameter Extraction
Question Ranking
s_1 s_2 s_3 s_{n_s}
t_1 t_2 NULL t_{n_s}
$(t_1, p_{t_1,1})$ $(t_1, p_{t_1,2})$ $(t_{n_s}, p_{t_{n_s},1})$
(t, p)
Job Posting
Job Sentences
Sentences w/ Template Label
Generated Questions
Top-K Questions

Identity Verification

Paper

Are You for Real? Detecting Identity Fraud via Dialogue Interactions

IDEA

The **fake** personal information of a loan applicant:

School	Company	Residence
Nanjing University	Baidu	No.3 Gulou Street

System: Which university did you graduate from?
Applicant: Nanjing University.
System: Which company do you work for?
Applicant: Baidu.
System: Where do you live?
Applicant: No.3 Gulou Street.
Decision: Non-Fraud (Wrong)

System: When was Nanjing University founded?
Applicant: I am not sure.
System: Which is the nearest subway station to Nanjing University?
Applicant: I am not quite clear.
Decision: Fraud (Correct)

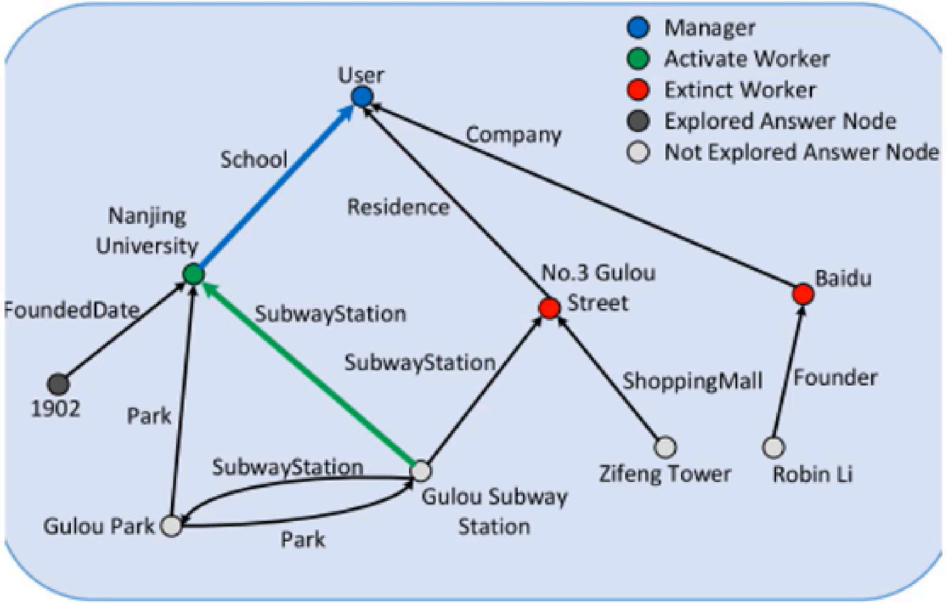

- knowledge graph-based questions that fraudsters are unlikely to know.

Human-AI Customer Service

Paper

Service Group: A Human-Machine Cooperation Solution for Group Chat Customer Service

Scenario

National Taiw

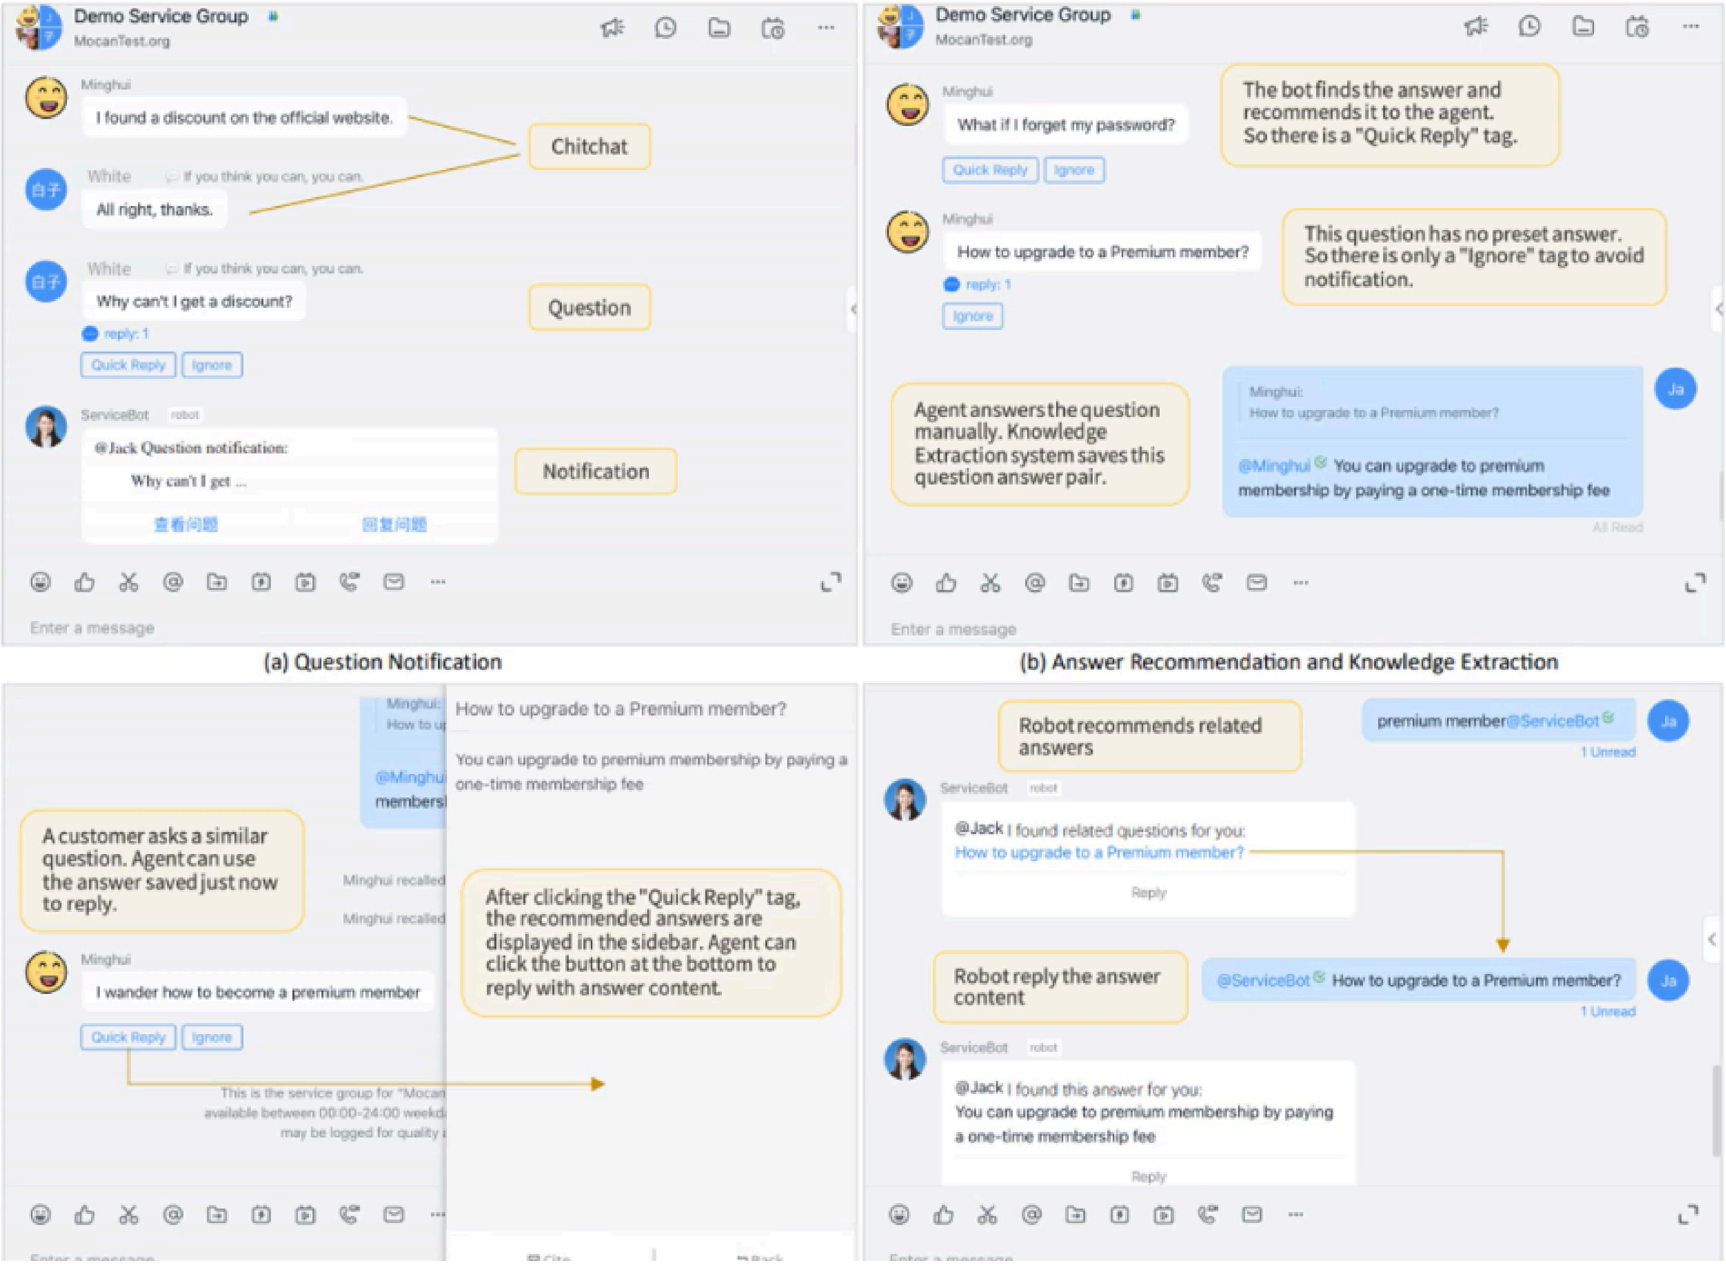

(a) Question Notification

(b) Answer Recommendation and Knowledge Extraction

(c) Answer Recommendation and Knowledge Extraction (d) Question Answering

-
- knowledge graph-based questions that fraudsters are unlikely to know.

Debt Collection Chatbot

papers :

Two-stage Behavior Cloning for Spoken Dialogue System in Debt Collection

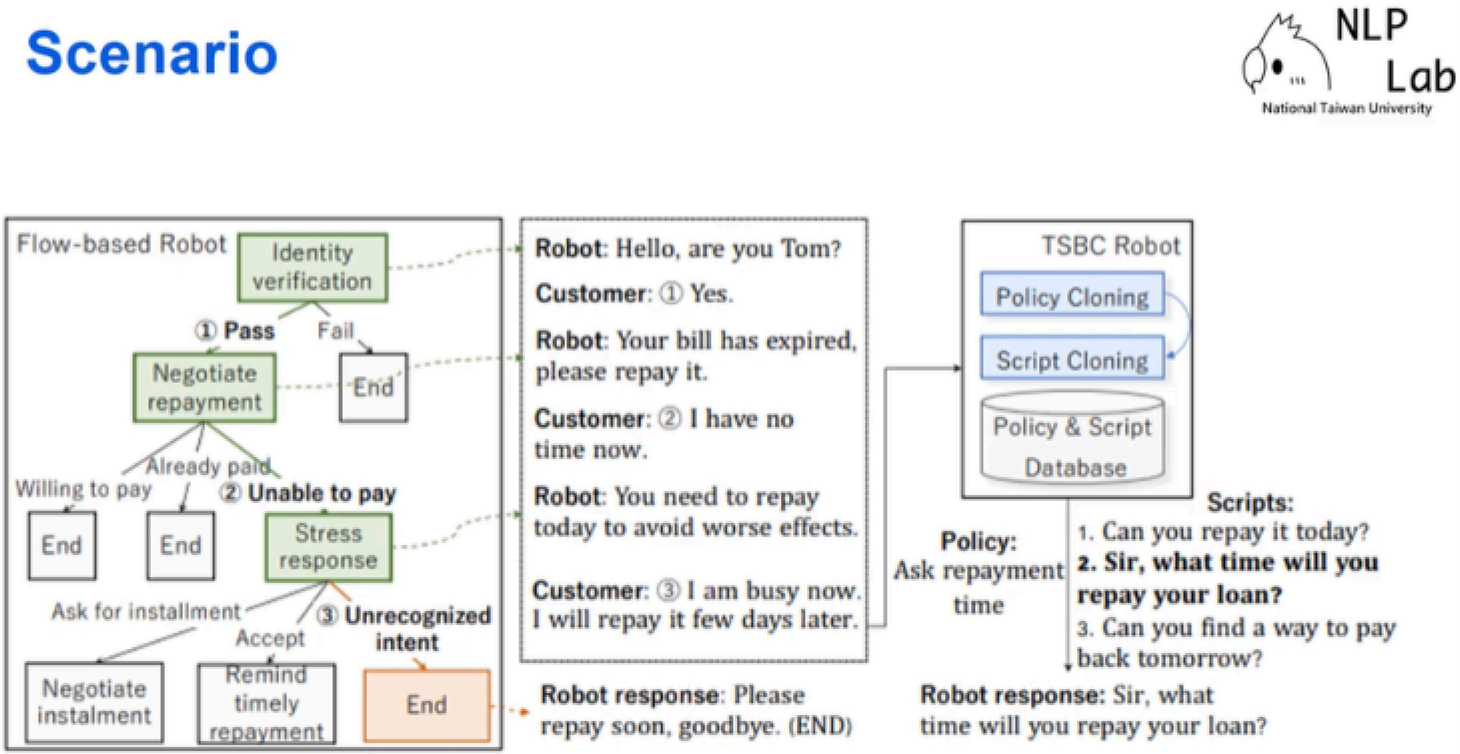

-

Debt Collection Chatbot

papers :

Unknown Intent Detection Using Gaussian Mixture Model with an Application to Zero-shot Intent Classification

- **TSBC** replaces fragile, hardcoded dialogue decision trees in debt collection with a **two-stage behavior cloning model** that learns high-level collection policies and selects candidate scripts directly from unannotated human collector call logs.

Learning with Weak Supervision for Email Intent Detection

Federated Learning

Financial institutions cannot freely share customer data. Instead of sending data,

- each institution trains locally.
- Only model parameters are shared.

papers :

Federated Meta-Learning for Fraudulent Credit Card Detection

Fraud Detection

papers :

Interpretable Multimodal Learning for Intelligent Regulation in Online Payment Systems

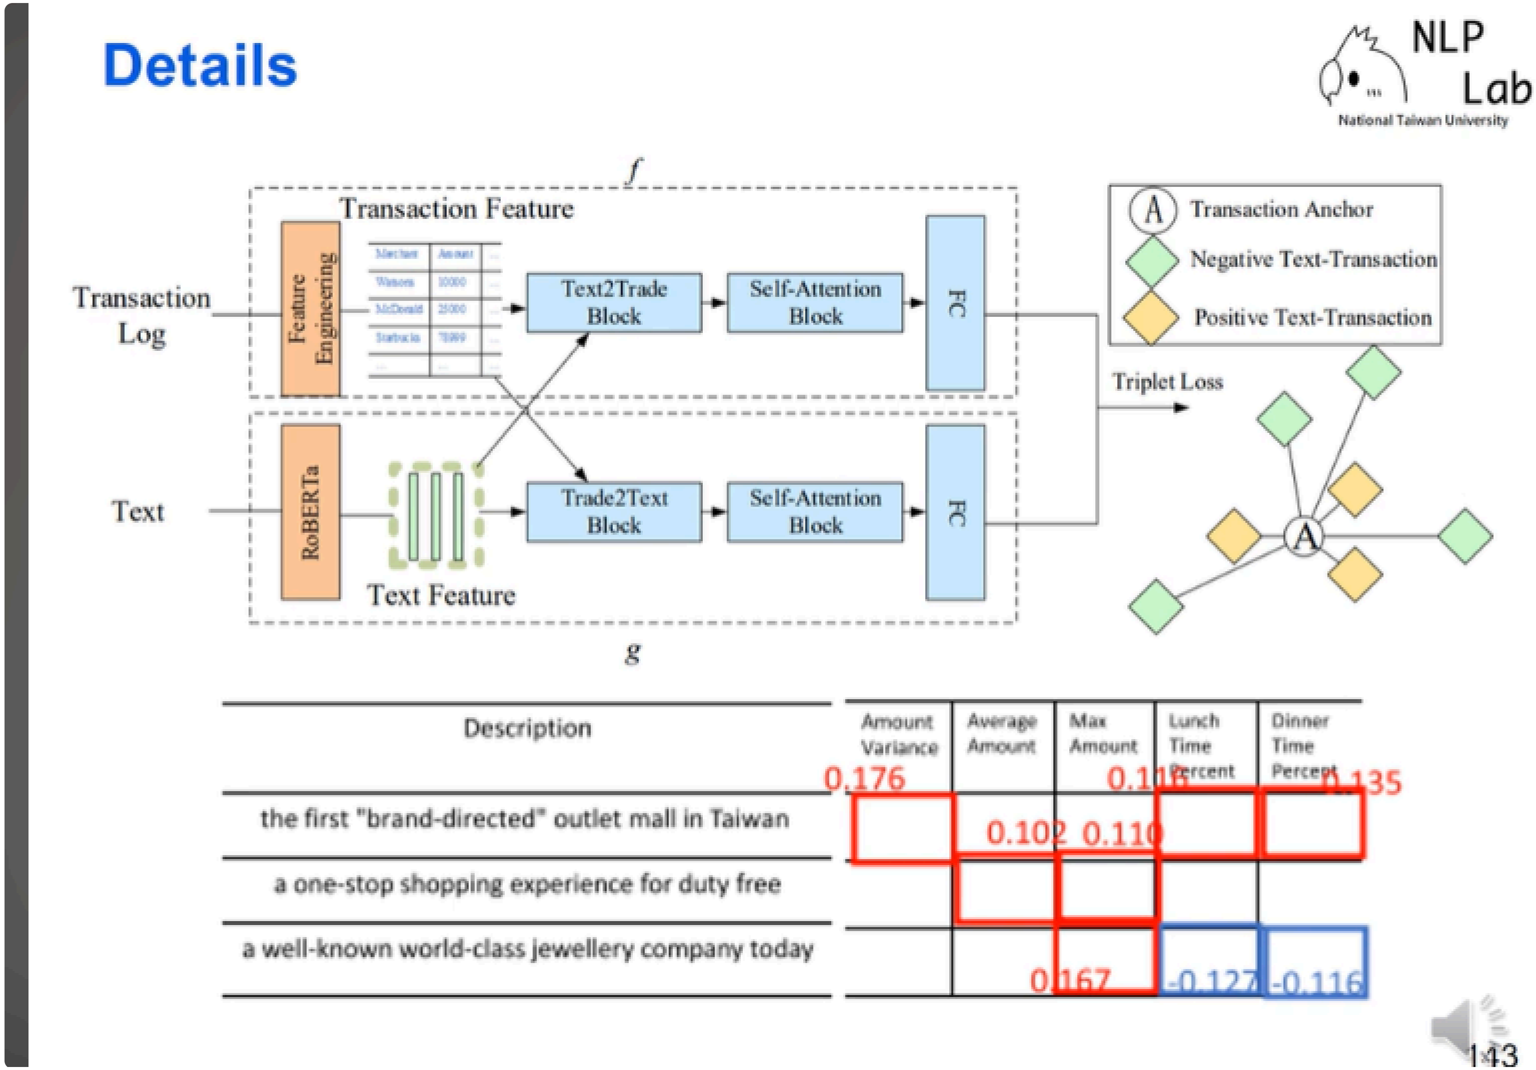

- built an AI detective called **CIAN** (used on Tencent's WeChat Pay) that checks if a shop's **written description** matches what they are **actually selling**.
- **CIAN** uses **cross-modal and intra-modal attention** to fuse merchant text descriptions with transaction flows, pairing it with a **low-rank matrix approximation explainer (CIAN-Explainer)** for interpretable regulatory risk detection in WeChat Pay.

The Behavioral Sign of Account Theft: Realizing Online Payment Fraud Alert

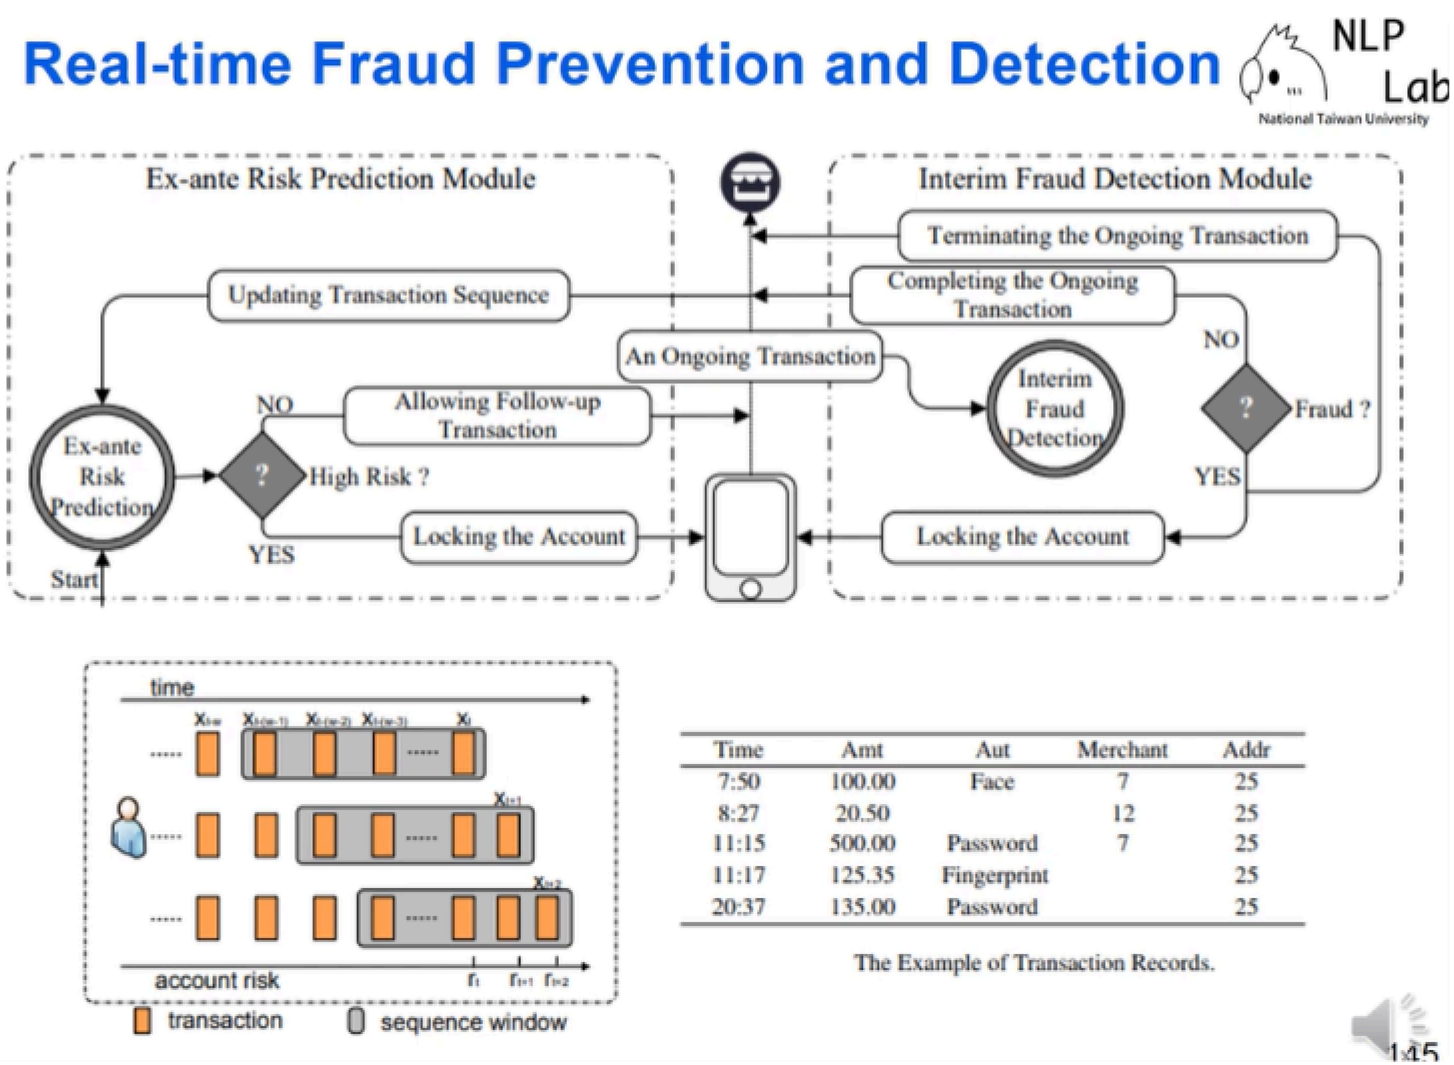

- This paper proposes an **ex-ante account risk prediction framework** that evaluates a user's recent transaction sequence window to predict **account compromise before a thief can initiate a fraudulent transfer**.

Numeracy in Financial NLP

Numbers carry much more meaning in finance than in ordinary text.

papers :

Learning to Generate Correct Numeric Values in News Headlines

- 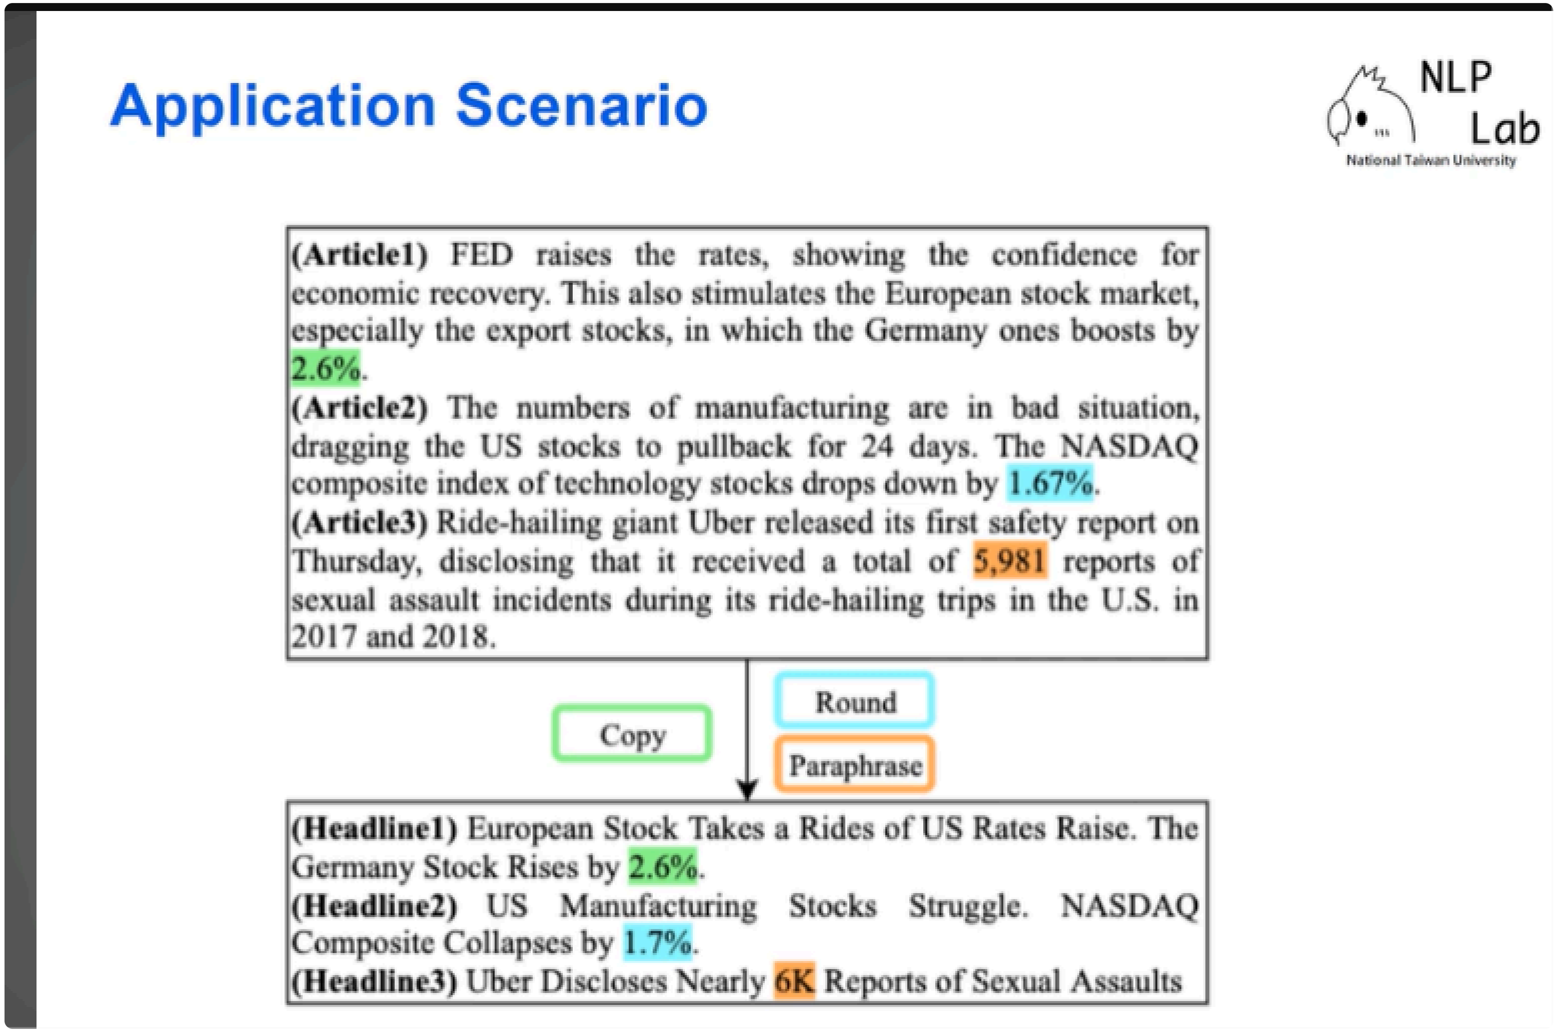

- Introduces a pretraining method to correctly select and generate numerical values when creating financial news headlines.

Table Understanding

Financial reports contain

- balance sheets
- income statements
- financial tables

These cannot be understood using plain text models.

---

papers :

TABERT: Pretraining for Joint Understanding of Textual and Tabular Data

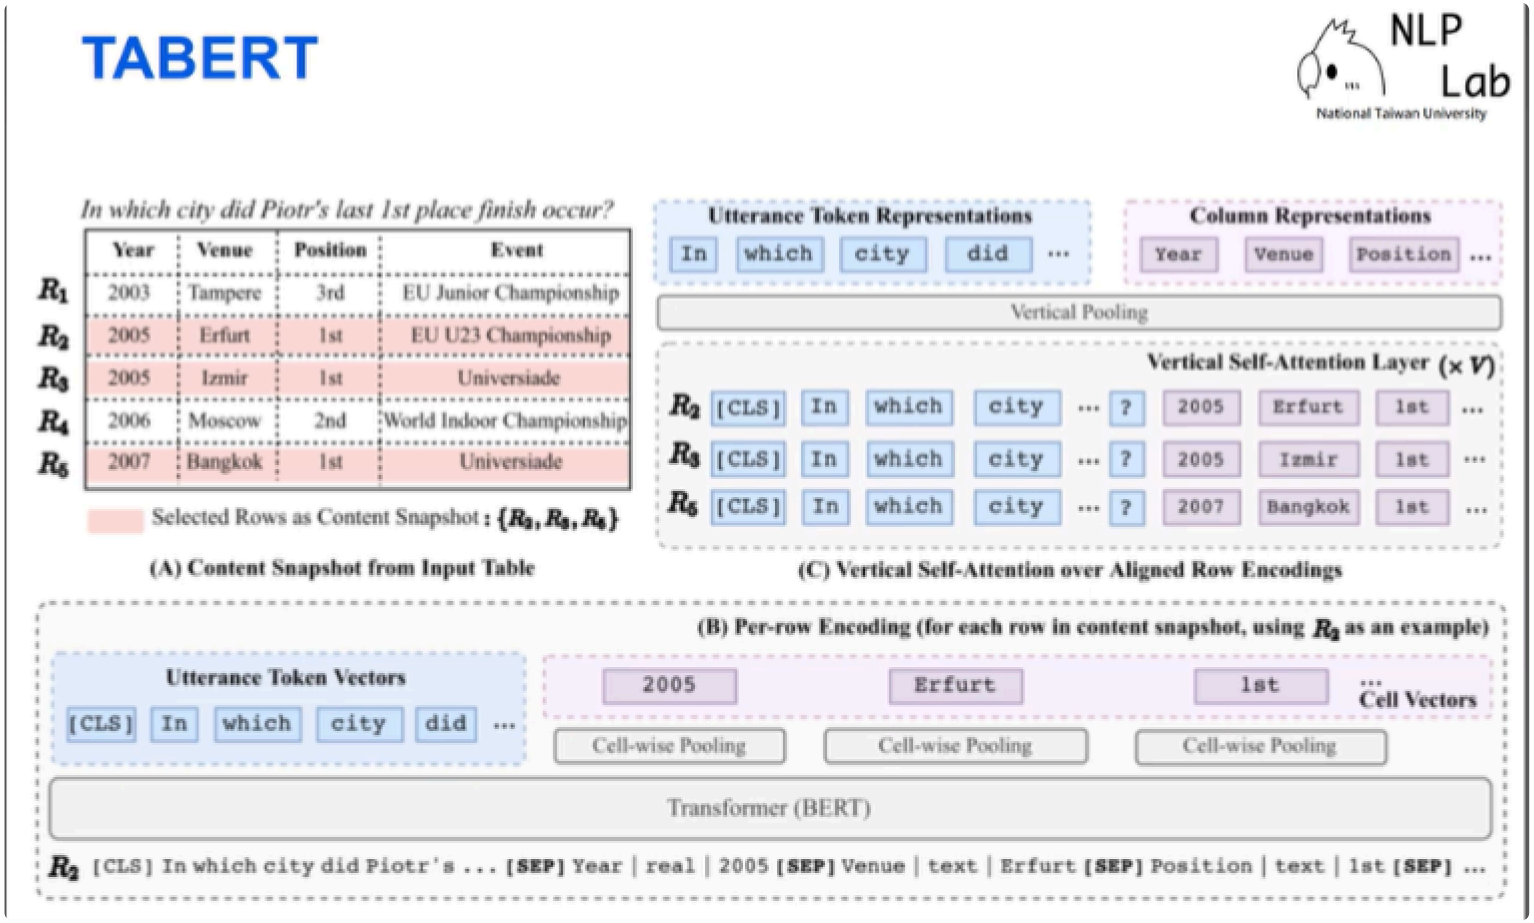

TAPAS: Weakly Supervised Table Parsing via Pre-training

SPot: A Tool for Identifying Operating Segments in Financial Tables

Problems Faced

Challenges – Dataset

1. Data cannot be shared
   - **Alternative Plans**
     - Release De-Identification Dataset
     - Using Social Media Data

2. Few benchmark datasets for backtesting

3. Data held by different financial institutions
   - **Return data to its owner**

Challenges – Task

1. Cross-Platform Analysis
2. Encoding Data with Temporal Information
3. **Explainable**
    - EU: Ethics Guidelines for Trustworthy AI
        - Traceability Mechanisms
        - Auditability
4. Lifelong Learning Algorithm
5. Automatic Factor Discovery
6. MultiModal
7. Personalized Assistant
    - **Investment Advisor**
    - Insurance Advisor
    - Customer Service

Model Limitations in Fine-Grained Tasks: While current models can detect large anomalies, they struggle with **"smaller modifications"** in financial numerals, which is a major issue for detecting false information

Regulatory Adaptation: Developing models that can automatically understand and check compliance when **regulations or laws change** is a difficult but essential task for future research

Conclusion

Take-Home Message

NLP Lab
National Taiwan University

- Traditional Financial Institution
  - **AI-assisted Operation**
    - Provide Better Service
      - Personal Finance → **Scenario-Based Financial Services**
    - Automatically KYC
      - Corporate Finance → **Credit and Risk Management with AI**
  - **Acquisitions and Partnerships**

- Academics & FinTech Startups
  - Novel Application Scenario & New Business Model
  - Tailor-Made Domain-Specific Models

- Papers in Finance and NLP
  - ACL Anthology: https://www.aclweb.org/anthology/
  - SSRN: https://papers.ssrn.com/sol3/DisplayJournalBrowse.cfm