A tool for Matching CAE Institutions' Student Skills to Job Requirements

Dr. Ram Dantu
Ram.Dantu@unt.edu
University of North Texas
Denton, TX

Tyler Parks
TylerParks@my.unt.edu
University of North Texas
Denton, TX

Leslie Delval
LeslieDelval@my.unt.edu
University of North Texas
Denton, TX

Thomas McCullough
ThomasMcCullough@my.unt.edu
University of North Texas
Denton, TX

Acknowledgements

- This Careers Preparation National Center product was funded by a National Centers of Academic Excellence in Cybersecurity grant (H98230-22-1-0329), which is part of the National Security Agency.

Introduction

- The mismatch between job seekers and employers is contributing to the increasing talent gap in Cybersecurity and Computer Science
- Prospective employees must understand how their skills align with current job market trends
- Mediums such as job postings and classroom assessments, provide primary sources of information regarding what employers are seeking

Introduction (cont.)

- CyberSeek's supply and demand heat map shows in the United States:
    - 700,000 unfilled cybersecurity positions
    - Out of 1.8 million total cybersecurity jobs nationwide (40% are unfilled)

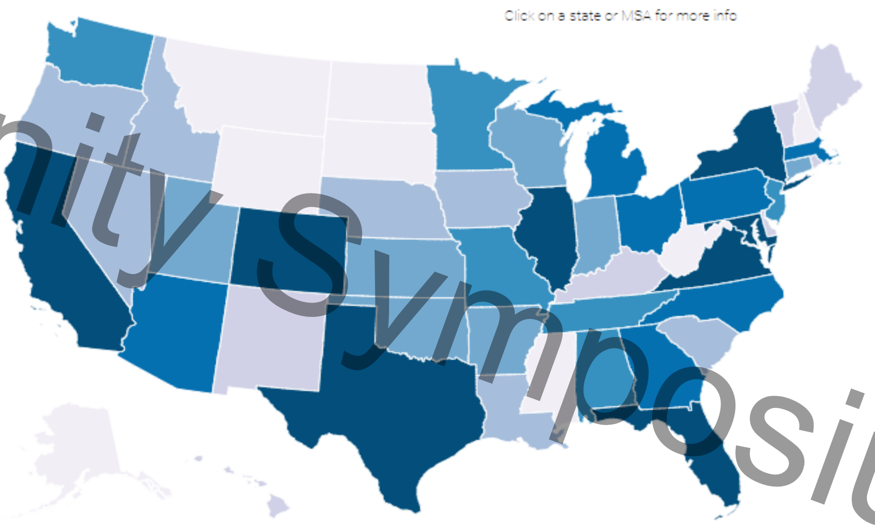

Overview

- This tool provides users with a platform to access a side-by-side comparison of classroom assessment and job posting requirements
- Uses techniques and methodologies from NLP, Machine Learning, Data Analysis, and Data Mining
- Analyzes job postings and classroom assessments, extracts and classifies skill units within, then compares sets of skills from different input volumes.
- Effectively provides a predicted alignment between academic and career sources, both federal and industrial

Overview (cont.)

- This tool has an overall accuracy score of **82%**, and an alignment score of only **75.5%** between the input assessments and overall job postings.
- These results are based on 50 UNT assessments and 5,000 industry and federal job postings examined
- Potential to be expanded into multiple subjects, across multiple job posting websites

Applications

- This tool is best suited for students who want to identify requirements from a job posting, classroom assessment, or potential resume
- Additionally, academic institutions could use this tool in conjunction with LMSs to generate a student's skill set
- Rubric creation and mapping is another potential feature that these institutions could implement using using the technologies mentioned in this tool

Motivation

- Learning outcomes are not only used to create academic material but also used by companies to detail their job postings.
- The requirements listed in assessments and job postings reflect the learning outcomes that a given party wishes to cultivate.
- No skill standardization:
  - Skill mismatches
  - Gaps in the hiring process
- Some standards:
  - Approaching to mend the gap

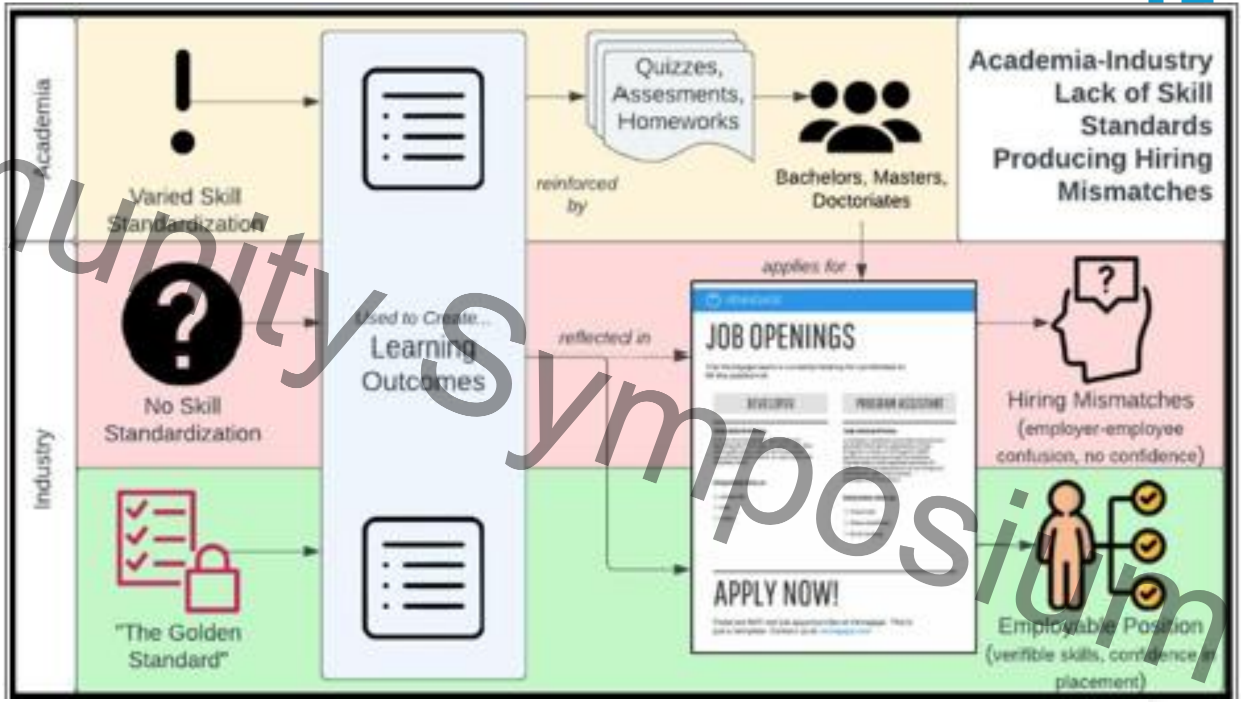

Problem

- There exists no widely-used "golden standard" set of skills, knowledge, or experience that is used by both academia and industry.
- Lack of transferability
  - Learned material from university
  - Applicable skills in-industry
- Without a standard of skill and knowledge items or a meaningful way to provide transparency between industry and academia, the entire hiring process succumbs to subjectivity

Our Solution

This tool proposes a combination of tools and methods to create a novel outcome alignment between given inputs.

- Filter out any job posting that is not an entry-level role
- Enabling job applicants' ability to accurately identify which positions their skills match closest to, will improve the frequency of correct job matches.
- Web scraper parameters are flexible; able to expand to other job posting websites:
  - Glassdoor
  - Monster
  - LinkedIn
- End-user will have the ability to form a selection into a grading rubric
  - Sculpt curriculums and assignments alike
  - As the job market changes, so will the skills
  - As the skills change, so will the rubrics

Understanding CPNC Tool

Collection and Compilation of Careers (CPNC)

- Analyzes **job postings** and **classroom assignments**
- Summarize and convey given job posting information with a **list of related knowledge and skill units**
- Aims to tailor a **list of outcomes** to a given classroom assignment or job posting
- Will enable users to **view summarized job postings** and **compare their own talents and skills one-to-one**

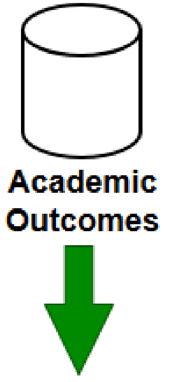

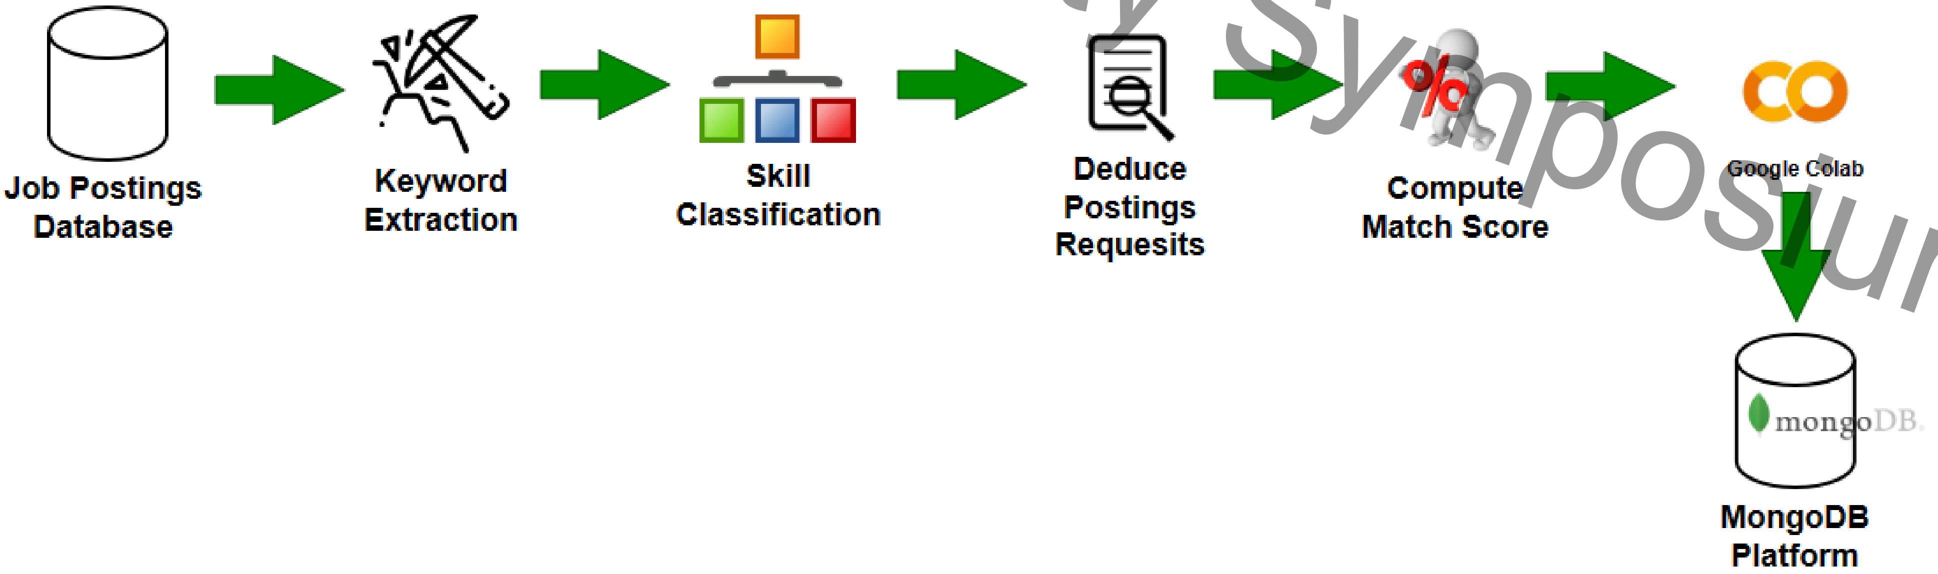

Inputs

This tool uses a variety of inputs across successive algorithmic steps. These inputs include:

1. Federal, Industrial job postings, and Classroom assessments *(from UNT, external assessments could also be used)*
    a. Then, using keyword extraction tools…
        i. Text volumes are processed
        ii. Only skill knowledge words and phrases remain

2. Knowledge unit collections and mappings from NIST/NICE
    a. A logistic regression model is trained
        i. Using ground truth skill standards and mappings
        ii. Enables the prediction of skill classifications

This classification creates learning outcomes from skill sets, where finally an alignment score between each input corpus can be calculated.

Overall Architecture Diagram

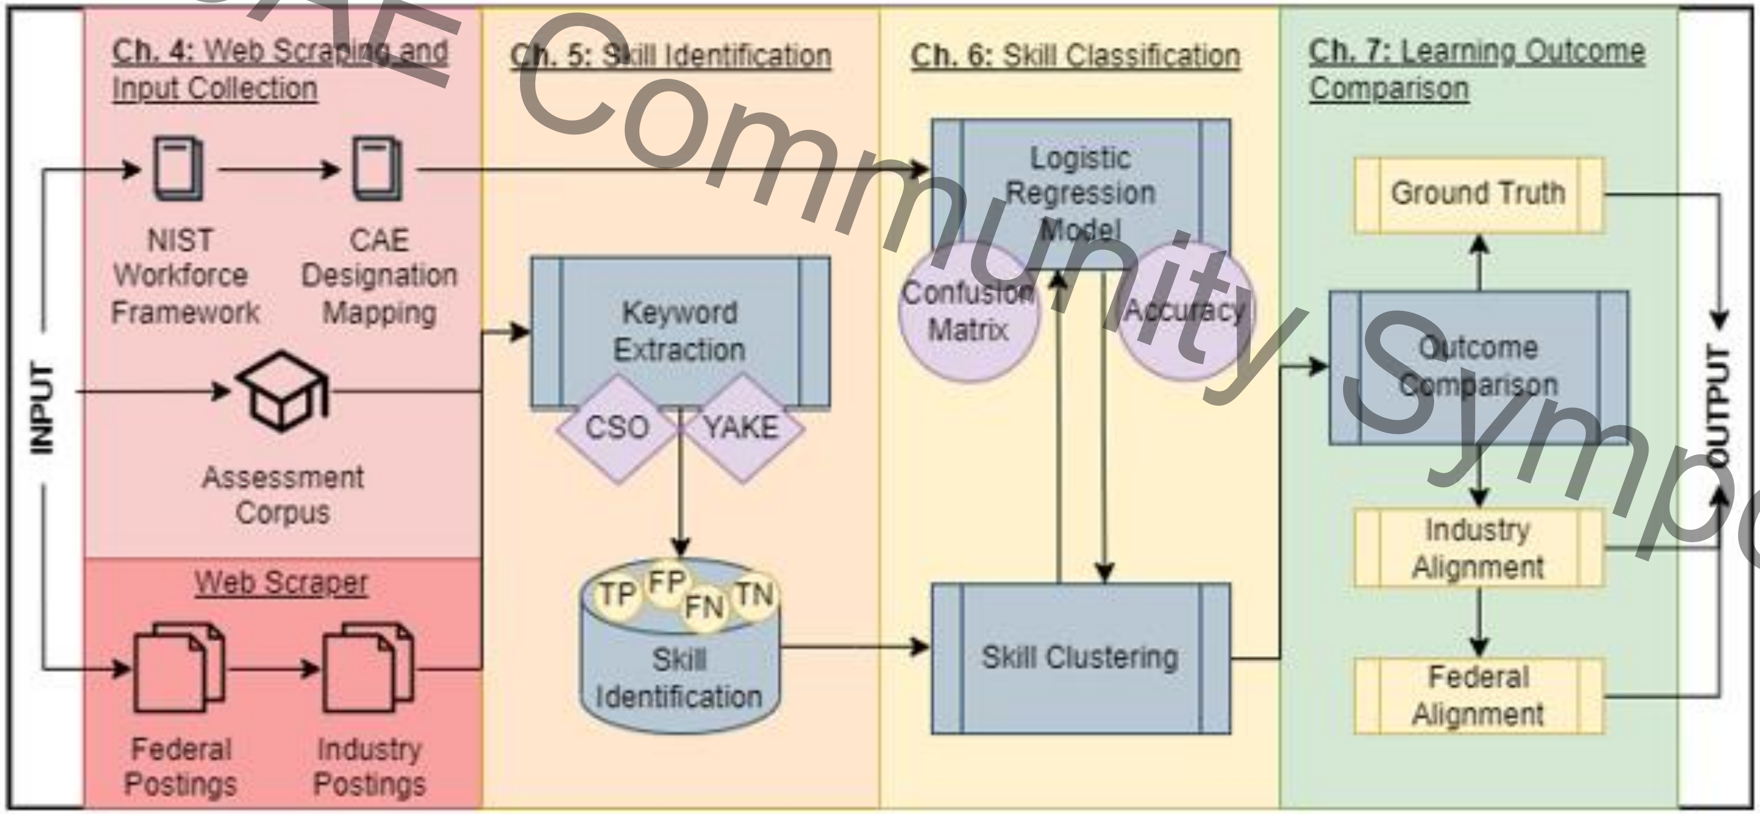

Notable Tools

- IDEs and other Interfaces
  - Python3
  - Google Collaboratory and Jupyter
  - Matplotlib
- Web Scraping
  - Custom scripts with:
    - requests
    - Beautifulsoup4
  - Octoparse
- Keyword Extraction
  - cso-classifier
  - YAKE
  - RAKE, skillNER (tested, not applied)
- String Classification
  - Sci-Kit Learn's
    - Logistic Regression Models
    - Confusion Matrices
    - Results Analytics
  - Labeled Skill Sets from Standardized Sources
    - NIST/NICE Workforce Framework
- Outcome Alignment
  - String Similarity
    - Word2vec
    - Fuzzy
    - spaCy

Results - Composite

Each value is color-coated to represent that section's value range from low to high, where:

- Red: 00.00 to 59.99
- Orange: 60.00 to 69.99
- Yellow: 70.00 to 79.99
- Green: 80.00 to 89.00
- Blue: 90.00 to 99.99

Outliers:

1. 96% **skill collection accuracy** on federal job postings
2. Individual **class accuracies**
   a. High of 92%
   b. Low of 47%
3. Low minimum value for all **outcome comparisons**
   a. 58.79% -> 62.14%

Chapter	Measure Description	Sub-measure	High Value (%)	Avg. Value (%)	Low Value (%)
4	4.1 Federal Job Posting Collection	USAJobs.gov, others	-	96.20	-
	4.2 Industry Skill Collection	Indeed.com, DICE.com	85.02	82.25	73.34
5	5.1 Keyword Extraction Tool Validation	CSO-Classifier	-	86.74	-
		YAKE	-	83.87	-
		CSO + YAKE	-	87.00	-
	5.2 Skill Extraction Accuracy	Classroom Assessments	-	87.94	-
		Federal Job Postings	-	78.15	-
		Industry Job Posting	-	82.50	-
6	6.1 Logistic Regression Model Prediction	-	-	76.00	-
	6.2 Training Data Validation	Most, Least Accurate Classes	92.00	76.00	47.00
7	7.1 Learning Outcome Alignment Score	Assessment vs. Ground Truth	77.85	68.78	58.79
		Assessment vs. Federal Job Postings	88.92	76.76	62.14
		Assessment vs. Industry Job Postings	86.48	75.02	59.14

Final Tool Performance

Put simply, the tool as a whole performs with an average accuracy of 82%.

The collection of classroom assessments used in outcome comparison has a:

- **76.76%** alignment with a collection of 200 federal job postings
- **75.02%** alignment with a collection of around 4,000 job postings
- **68.78%** alignment with a collection of ground truth skills.

These results are reflective of a moderate alignment (on conveyed requirements) between academia and industry, federal job postings.

These results can be decomposed to find that courses can be aligned as well, below to federal outcomes:

- CSCE5560: **62% - 88%**
- CSCE3550: **65% - 86%**
- CSCE5550: **68% - 87%**

Getting Started

- Go to tylerjparks.github.io or **scan** the QR code
- Try the tool using a job posting from our database

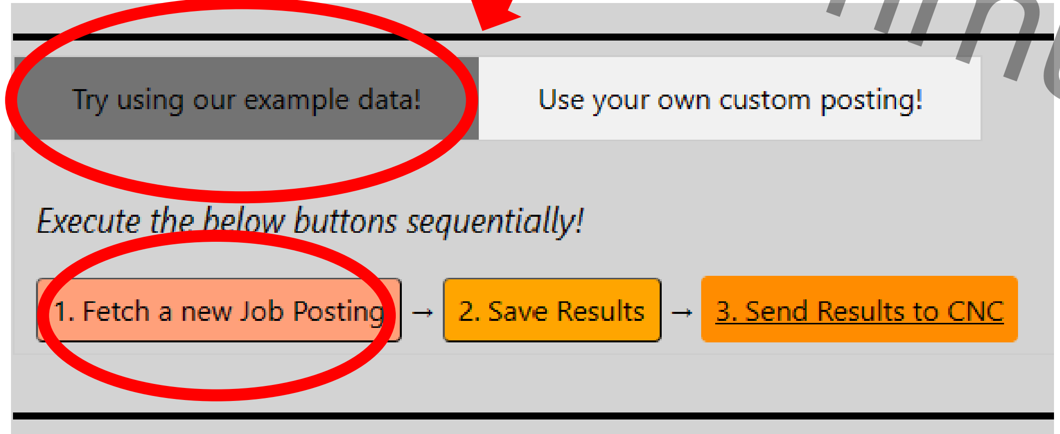

- or **copy-paste** the job posting of your choice

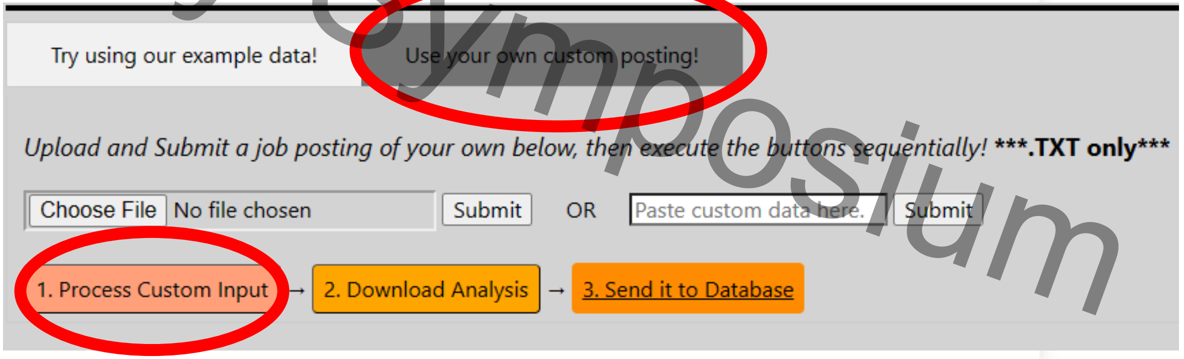

Getting Started (cont.)

- **Match score** of the job posting skills to your academic skills based on academic assessments

Step 1: Extract Keywords

List of Extracted Keywords with Scores

| 99.36%: security
| 99.15%: information security
| 98.34%: information
| 98.29%: ctds information security
| 98.06%: improve security operations
| 97.86%: threat
| 97.81%: threat intelligence
| 97.78%: security operations
| 97.71%: translational data science
| 97.36%: security events
| 97.15%: information security analyst
| 96.84%: information security officer
| 96.32%: information security tools
| 96.15%: imple information security
| 95.85%: data
| 95.31%: improve security
| 95.22%: security incidents
| 94.95%: bsd ctd
| 94.76%: wider information security
| 94.6%: intelligence

Keyword Average: 96.94%

Step 2: Generate Classification Groups

Classifications Used with Training Score

| 62.0%: Data Security
| 44.0%: IT Management/Services
| 52.0%: Cyber Threats
| 72.0%: Data Mining
| 45.0%: Cybersecurity Foundations
| 47.0%: Forensics

Model Average: 60.23%

Step 3: Categorizing Posting Requirements

Assigning Keywords to Appropriate Classes

| Data Security
| | security
| | information security
| | ctds information security
| | improve security operations
| | security operations
| | security events
| | information security analyst
| | information security officer
| | information security tools
| | imple information security
| | improve security
| | security incidents
| | wider information security

| IT Management/Services
| | information

| Cyber Threats
| | threat
| | threat intelligence

Step 4: CAE Match

| 90.3% match
| 100%: Data Security (item 1)
| 100%: Data Security (item 2)
| 100%: IT Management/Services (item 3)
| 70%: IT Management/Services (item 4)
| 100%: IT Management/Services (item 5)
| 61%: IT Management/Services (item 6)
| 100%: IT Management/Services (item 7)
| 100%: IT Management/Services (item 8)
| 100%: IT Management/Services (item 9)
| 100%: IT Management/Services (item 10)
| 70%: IT Management/Services (item 11)
| 100%: IT Management/Services (item 12)
| 61%: IT Management/Services (item 13)
| 100%: IT Management/Services (item 14)
| 100%: IT Management/Services (item 15)
| 100%: IT Management/Services (item 16)
| 72%: Data Mining (item 17)
| 100%: Data Mining (item 18)
| 72%: Data Mining (item 19)
| 100%: Data Mining (item 20)

New Technology

- Leveraging LLMs such as **Chat-GPT** we can
  - Increase accuracy
  - Adapt our Cybersecurity compatibility score to other domains
  - Improve access to our database with integrated chat functionality

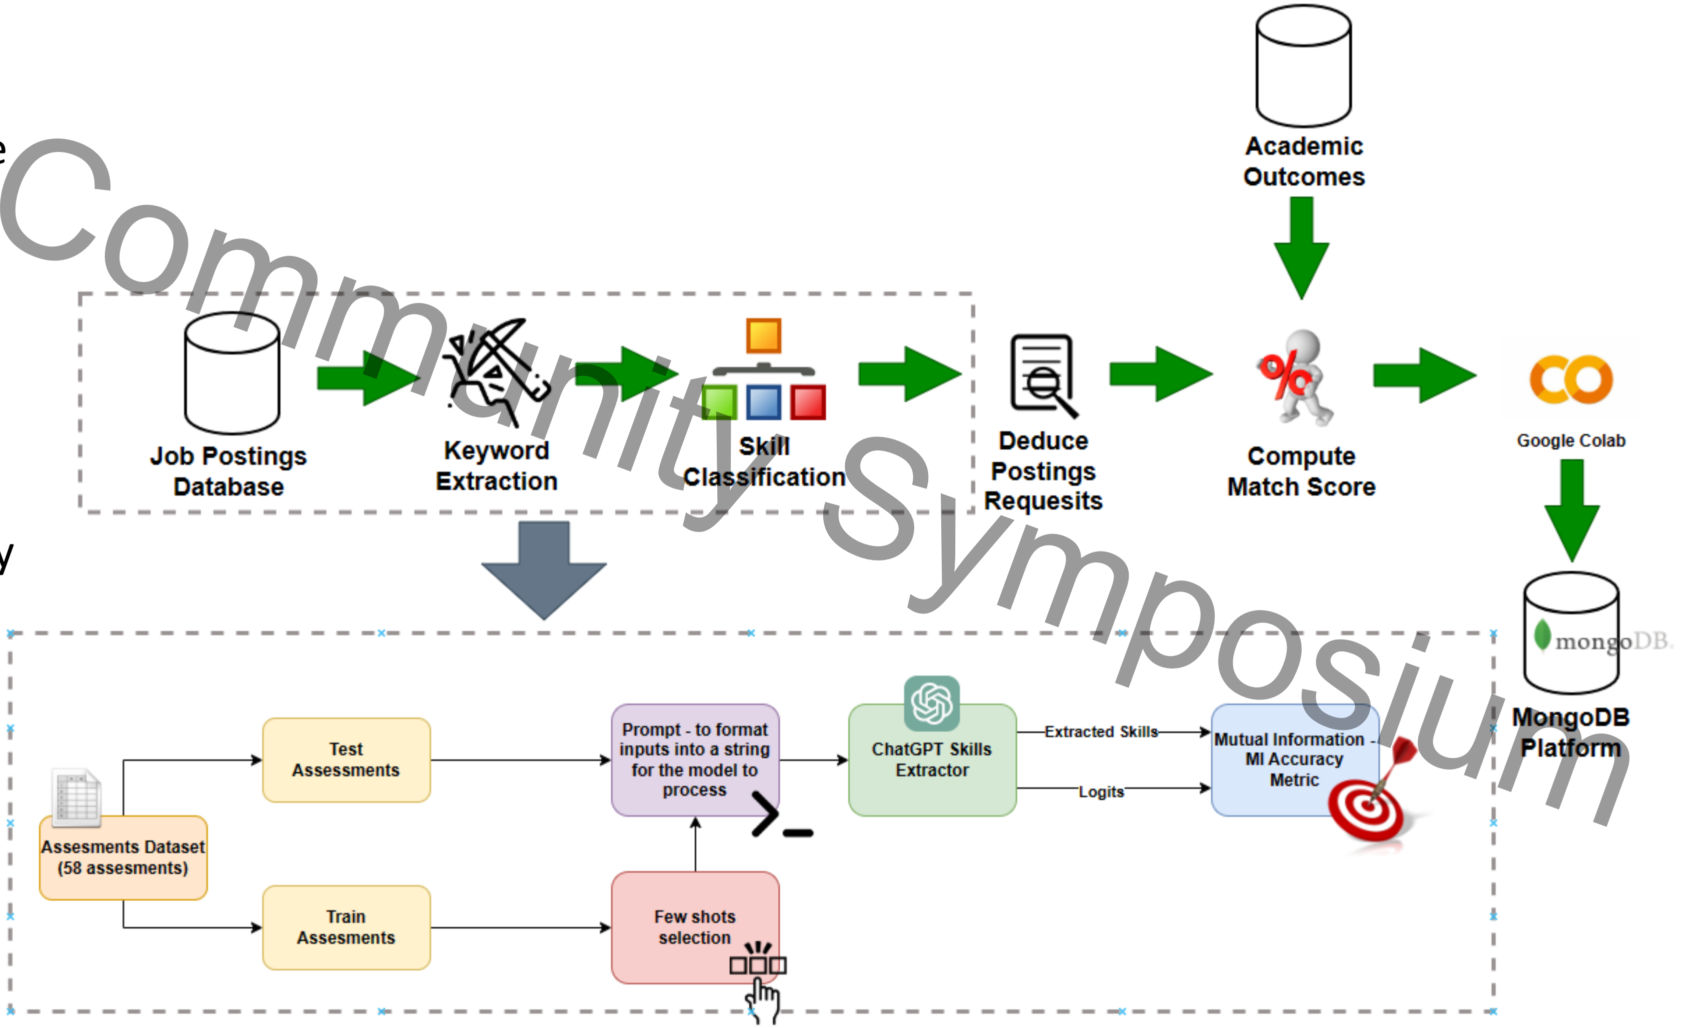

Skill Extractor Architecture

- Significantly **reduced hallucinations** through overconfident discounting and maximizing context length
- Variations reduced through **optimal few-shot ordering**
- In context learning can augment or **replace more expensive training** method

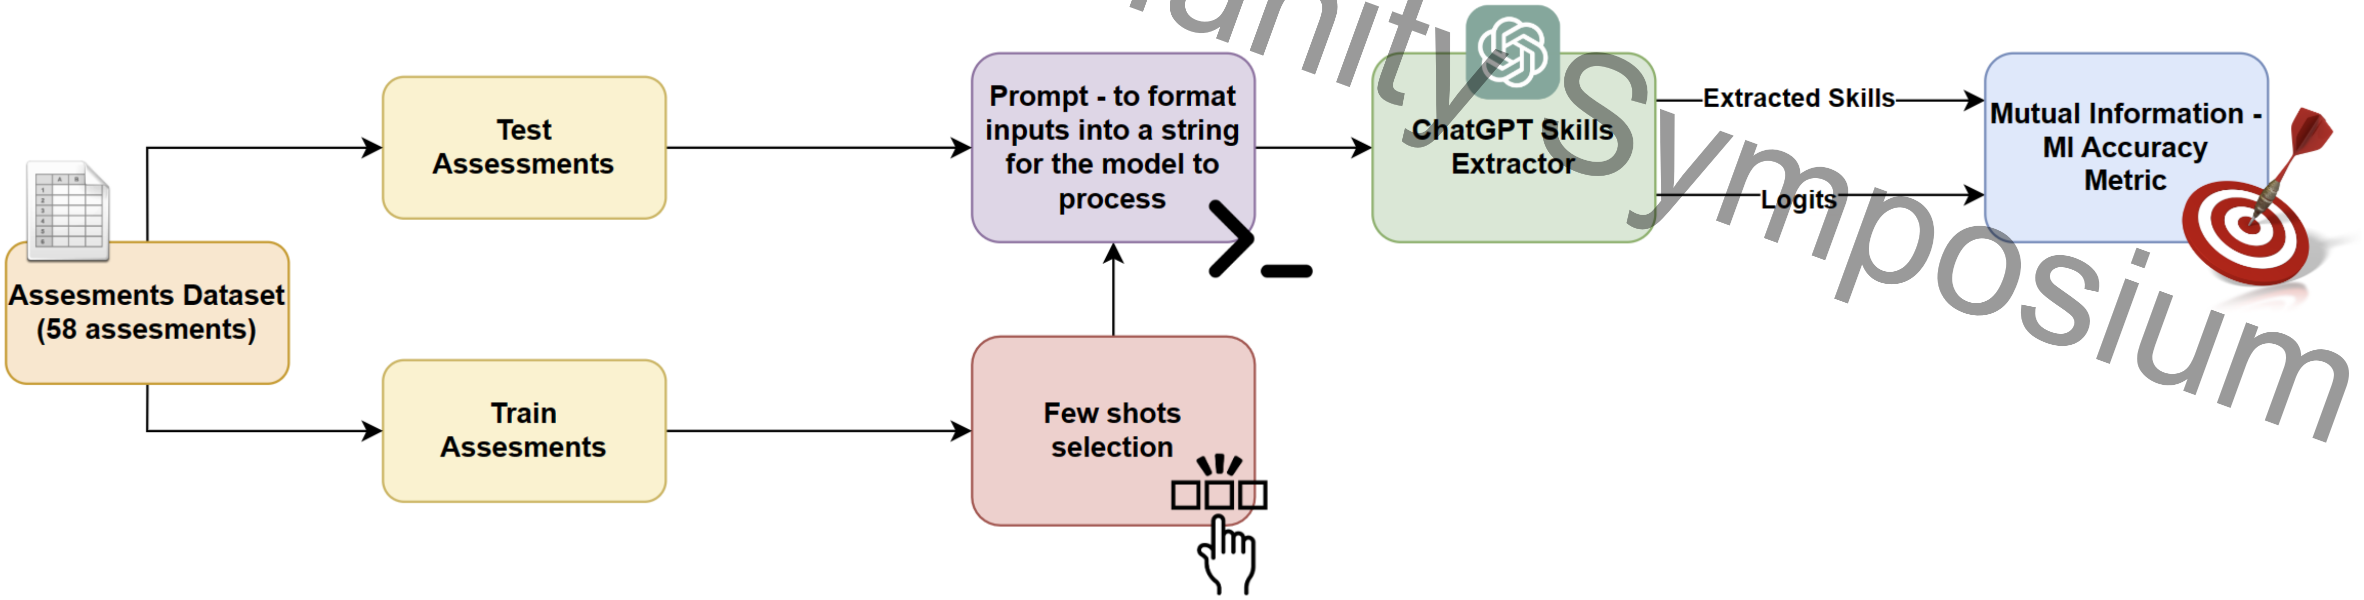

Skill Extractor Example Size Accuracy

- In this graph, we show that **as we increase the size of assessment segments, we increase the accuracy of extracted skills**. TP shows correctly extracted skills, while FP represents hallucinated skills. This is significant, as it indicates that **we should aim for the largest segments possible for the sake of skill extraction with Chat-GPT**

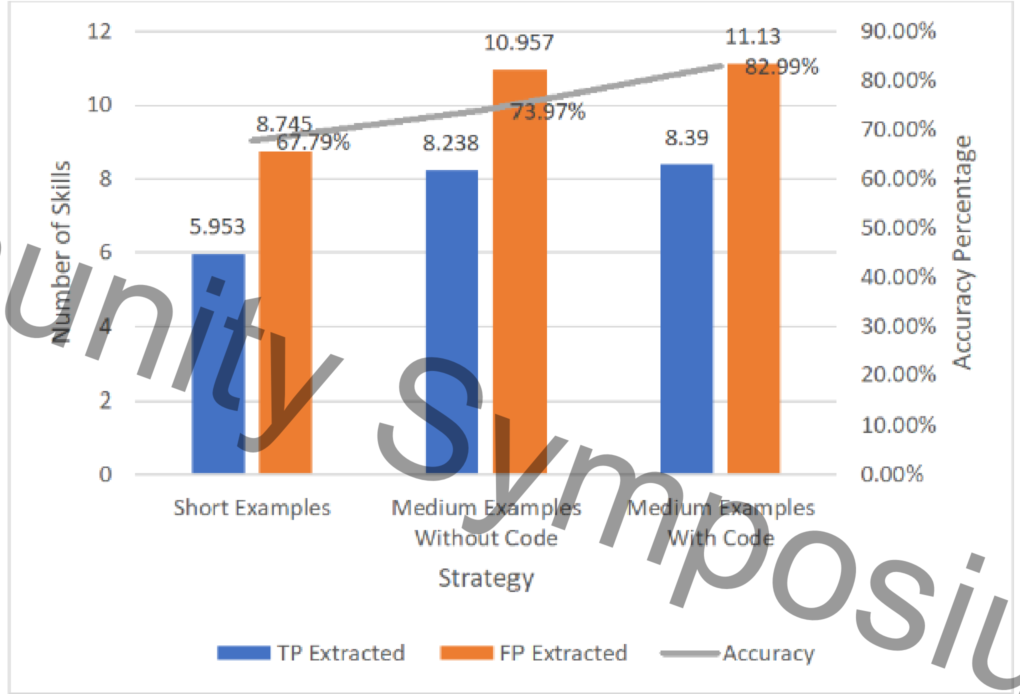

Figure 2 Table 3 Visualization. In this graph, you can notice that as we use more aggressive pruning strategies, the number of distinct skills goes down, while the accuracy of selected skills goes up. We will use this correlation to select for more accurate examples in the next section.

Removal of Overconfidence Strategy TP FP Accuracy

This table shows a correlation between overconfidence and hallucinations, which we observed with mutual information and hallucinations. While mutual information usually correlates with accuracy, **hallucinated skills made up the top 10-20% of extracted skills when sorting by mutual information**. By **pruning these hallucinations, we can avoid a major accuracy pitfall**.

Strategy	Accuracy	TP Num Skills	FP Num Skills	TP MI	FP MI
0% Top Overconfidence Removal	81.658%	5.705	7.965	1.355	1.324
10% Top Overconfidence Removal	87.830%	5.017	5.594	1.277	1.234
20% Top Overconfidence Removal	90.841%	4.921	4.959	1.211	1.213

Table 4

Rigor Extractor Architecture

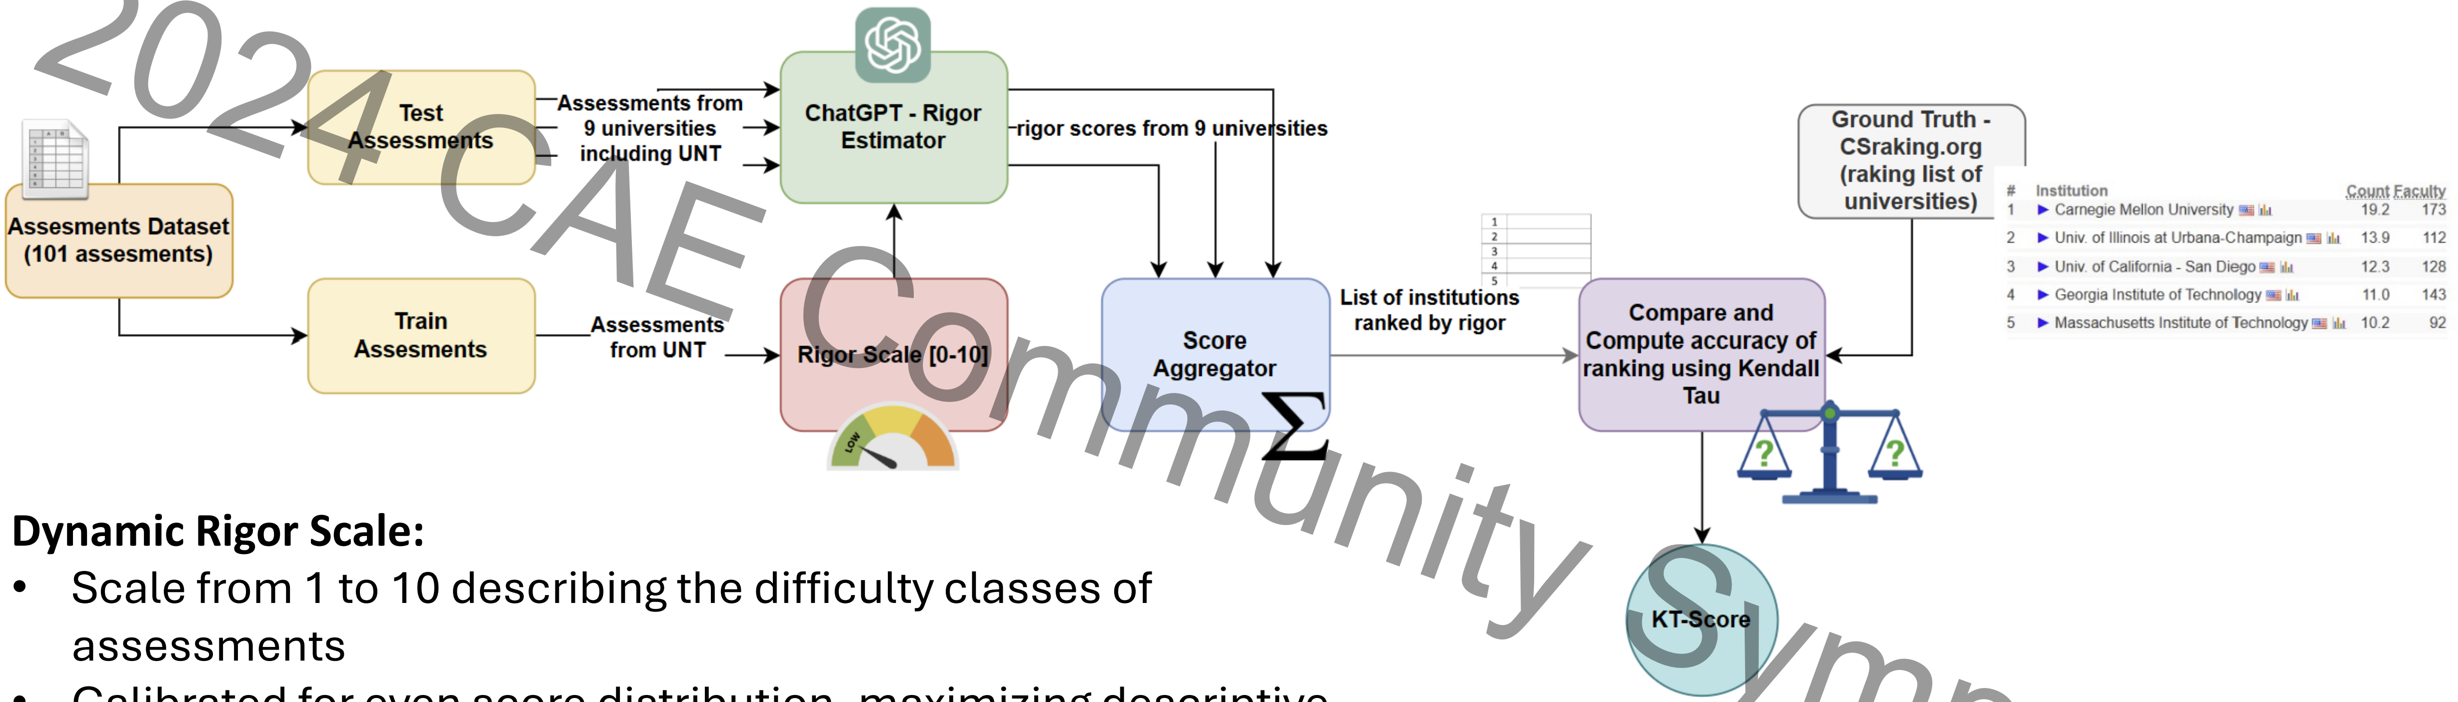

Dynamic Rigor Scale:

- Scale from 1 to 10 describing the difficulty classes of assessments
- Calibrated for even score distribution, maximizing descriptive power of the scale

Rigor Estimator:

- Single scores between 1 and 10 for each assessment
- Utilizes two new prompt engineering techniques for maximum accuracy
- Objective perspective, not influenced by student grades or personal instructor investment

Score Aggregator:

- Provides a single aggregate score for a course based on the extracted rigor across all assessments
- Random forest regressor trained on engineered mathematical features
- Significant correlation(.59) with ground truth rigor at csrankings.org

Conclusion

As this work has the potential be expanded into:

- Multiple fields
- Across multiple job posting websites
- Using multiple supporting tools

Most future work should be dedicated to optimization and improvement of current methods. Specifically, more work needs to be done on improving skill clustering and outcome alignment. These methods, in general, perform the worst of the 4 objectives:

- Aiming for upwards of 90% or more
  - Veritably be able to cover the breadth of cybersecurity and CS knowledge
  - Outcome alignments will have even more weight and meaning

Being able to verify and reproduce each step of this research is key. Additional work needs to be done to ensure that each step is empirically stable so that measurements can be reproduced with a degree of accuracy.

Thank you