Lecture2notes: Summarizing Lecture Videos by Classifying Slides and Analyzing Text

Hayden Housen
Natural Language Processing • Computer Vision • Machine Learning

https://haydenhousen.com
hayden@haydenhousen.com
GitHub: @HHousen
Mentor: Dhiraj Joshi @ IMB

Goal/Focus

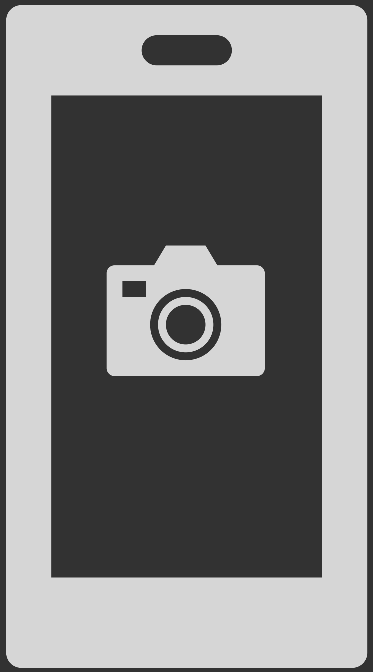

Phone

Presentation

Goal/Focus

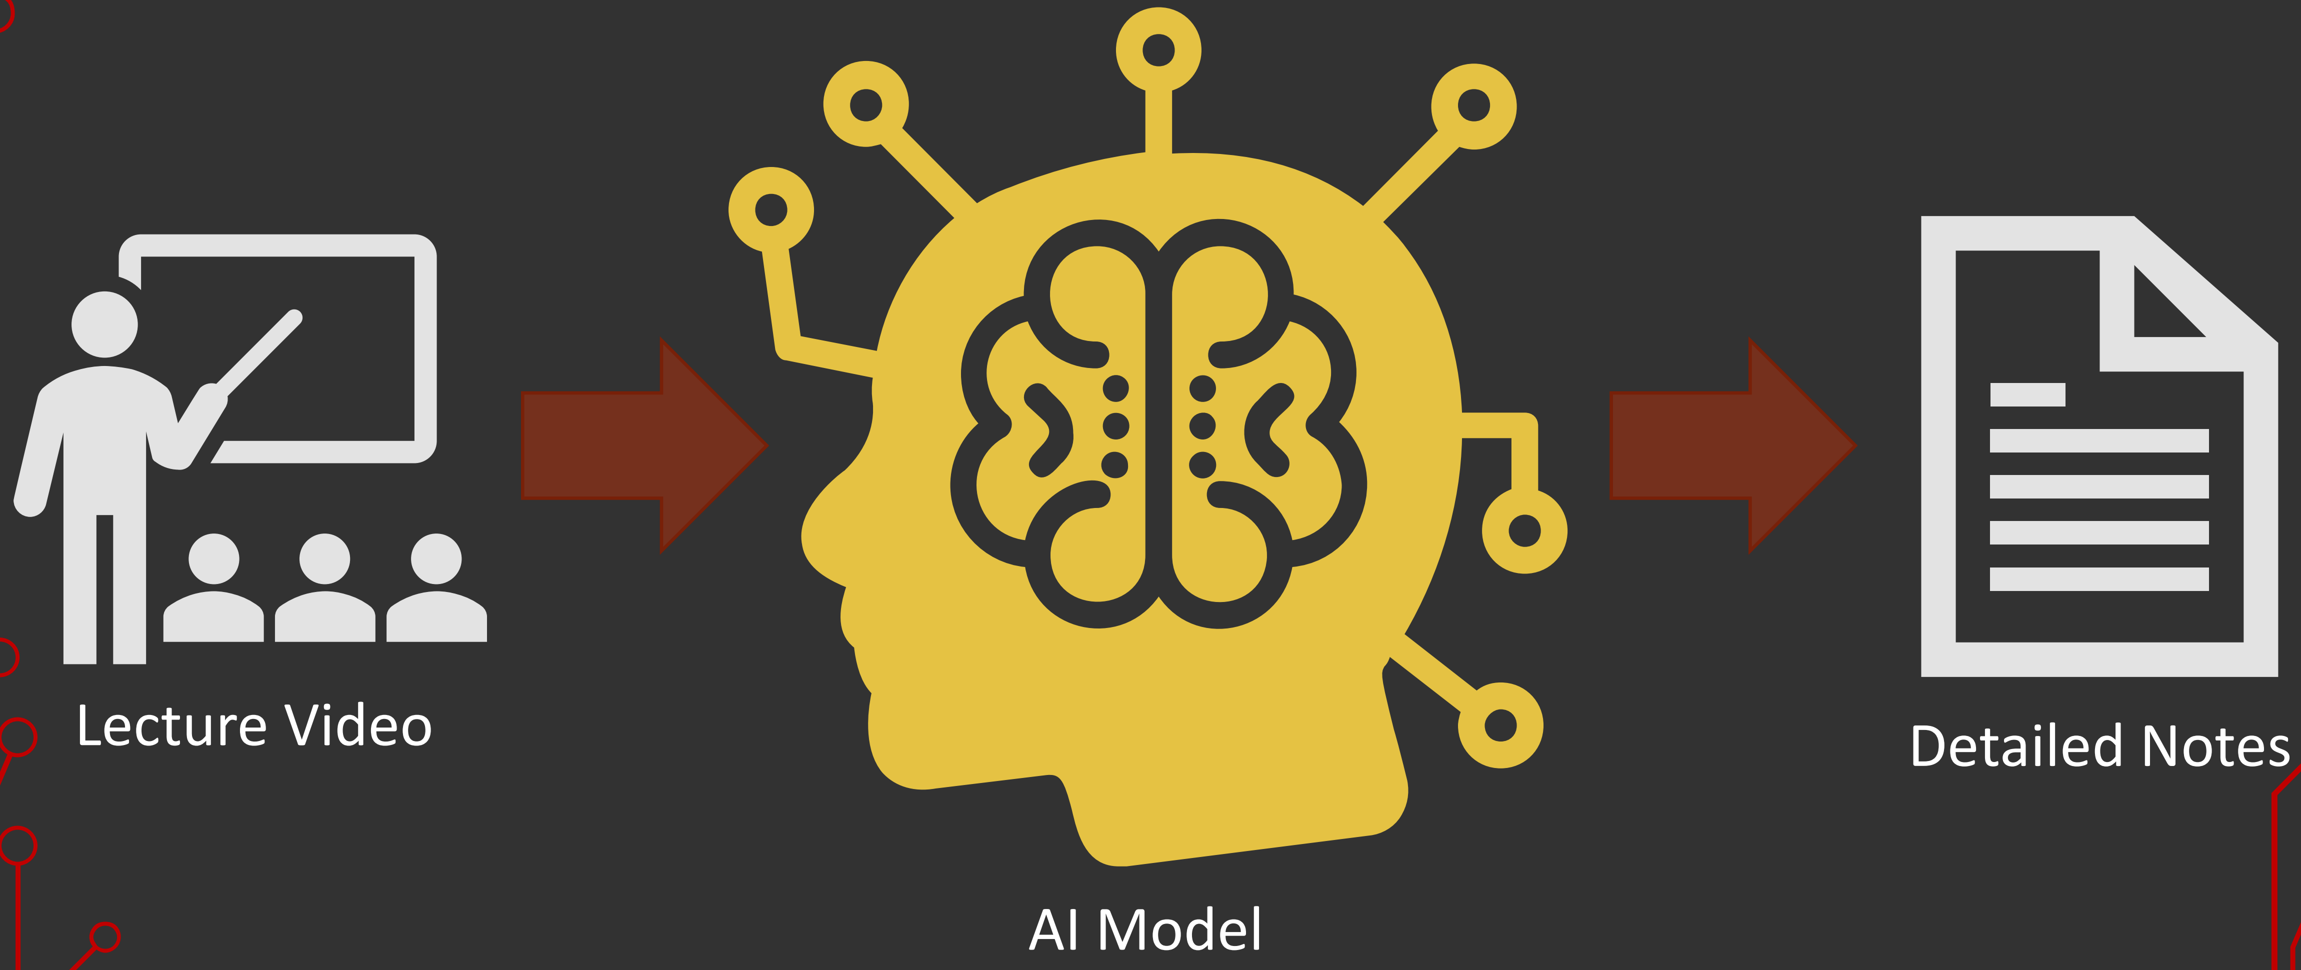

Gap/Hypothesis

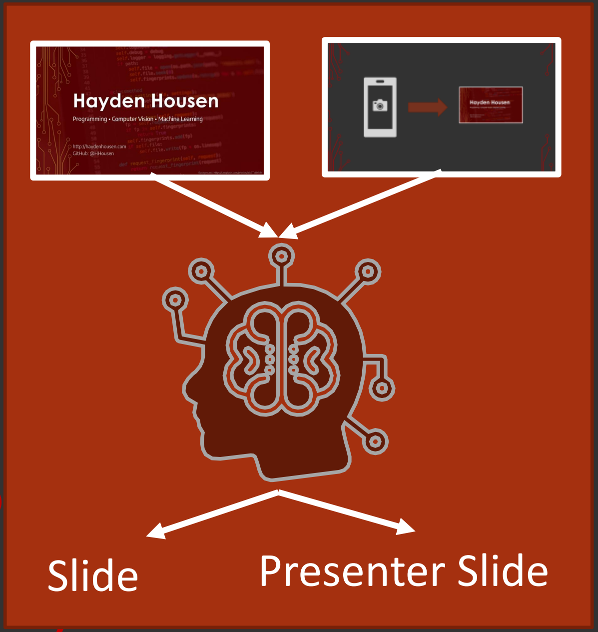

Slide Classifier

AI Summarization Models

End-to-End Process

Website

Brief Methodology

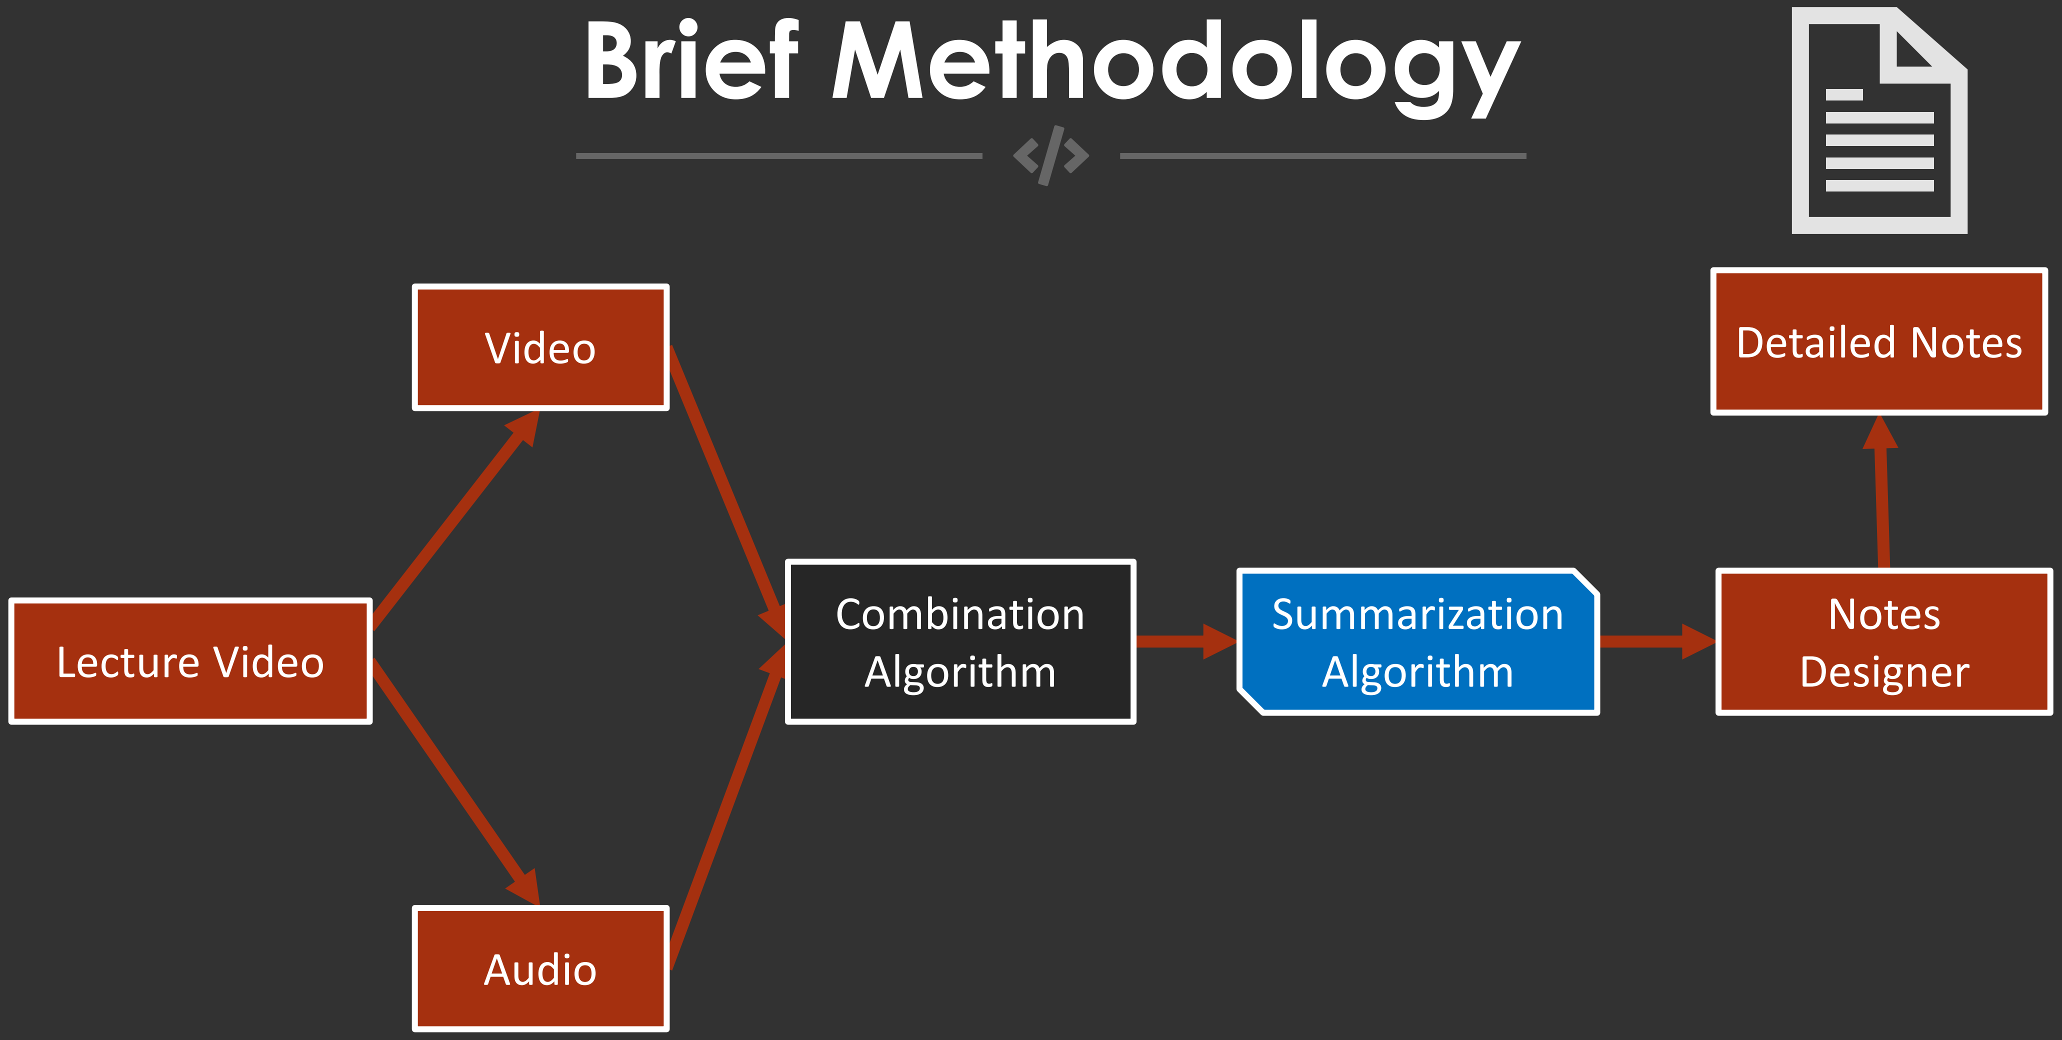

Computer Vision

“Teaching Computers To See”

- Face Recognition
- Emotion Analysis
- Crowd Analytics
- Sports: draw lines on field & highlights
- Medical Imaging
- Movie Special Effects
- Self Driving Cars

Importance

Humans

- **Vision** = vital for basic tasks
- Humans generate massive amounts of video (**Video data accounted for 75 percent of the total internet traffic** in 2017) [23]

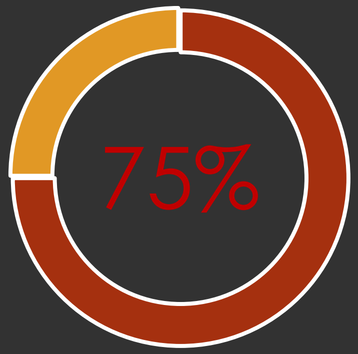

Computers

- For these reasons, computer scientists try to give computers sight for **half century** [1,2]
- **CV Goal** – teach computers how to understand digital images & videos at a high level

CV Basic Tasks

[14]

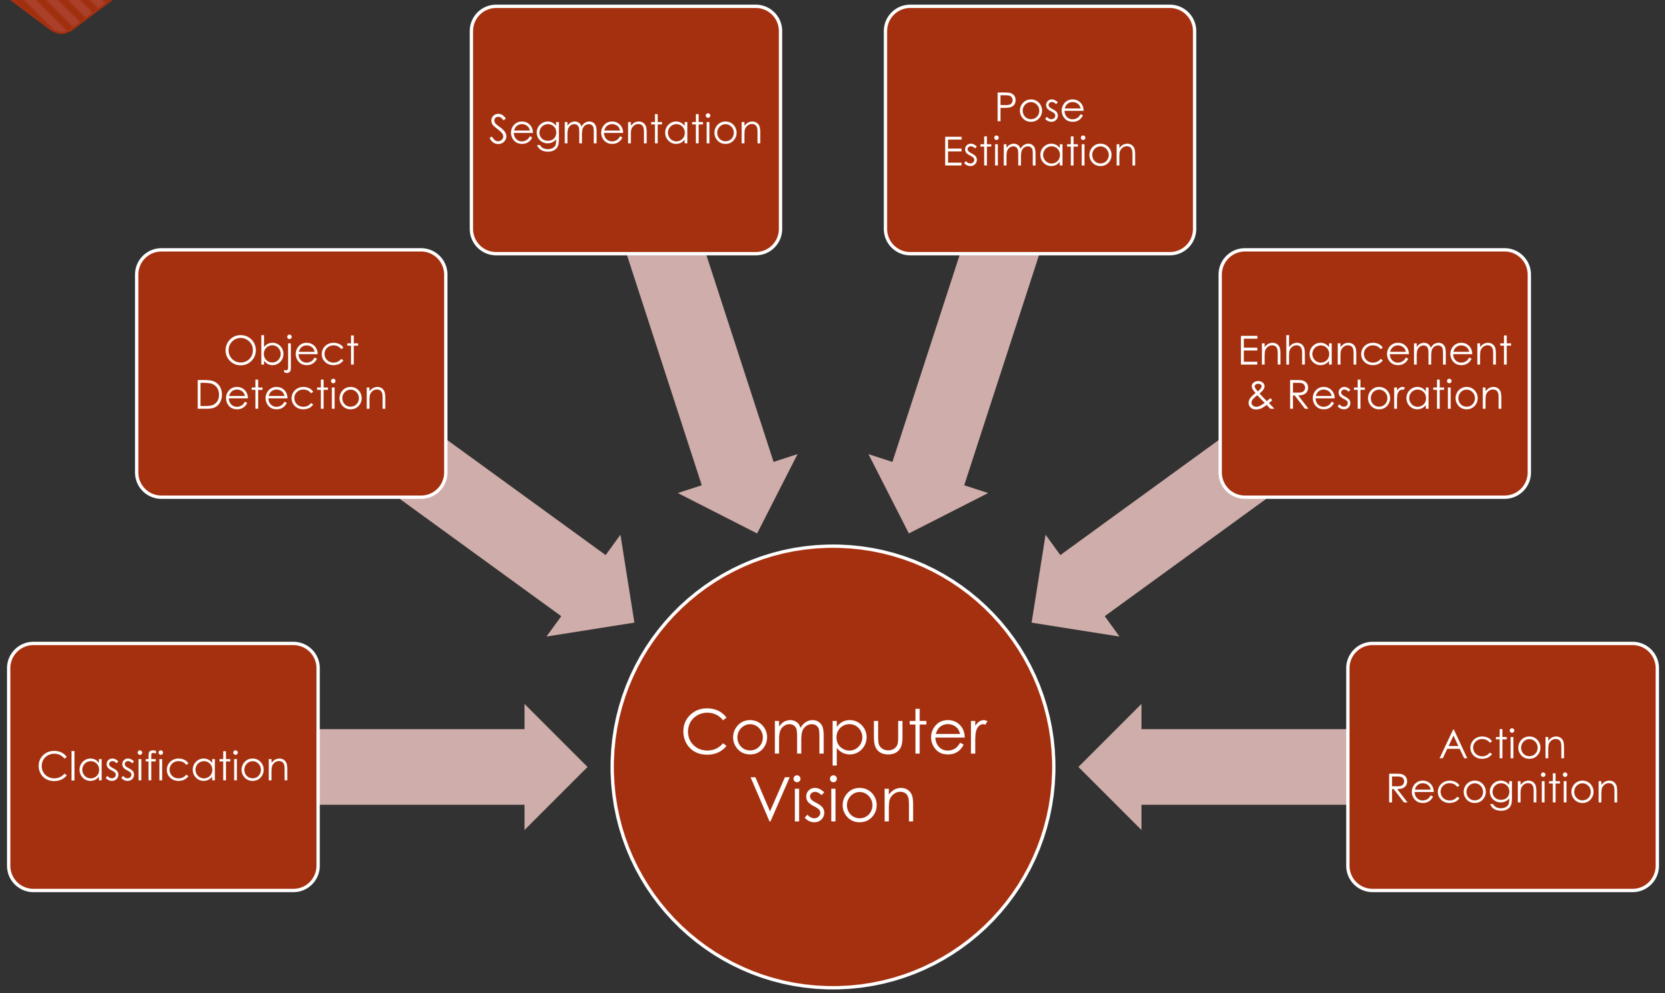

AI & Machine Learning

Artificial Intelligence

- Human intelligence exhibited by machines [3]

Machine Learning

- Statistical techniques that enable machines to improve at tasks without being explicity programmed [4]

Deep Learning

- Multilayered neural networks learn from vast amounts of data

Parse data

↓

Learn from it

↓

Predict something in world

Note Taking – Literature Review

- **Note taking** is almost a **universal activity** among students.[27]
- **Students' notes** are generally **incomplete**, and thus not adequate for reviewing the material.[28]
- Those that review **instructor provided** notes **score higher** than those who review their own notes.[28]
- Students **prefer guided notes**[29] and course final exam **performance higher** for guided notes.[30]

Estimated Effects

- **Previewing:** [31] suggests that summarized lecture slides reduce the amount of previewing time required without impacting quiz scores
- **MOOCs:** Summaries enable quick skimming of the main points.
- Automated summaries will
  - Decrease time spent creating notes
  - Increase quiz scores (content knowledge)
  - Enable faster learning

Main Components

Group / Hypothesis

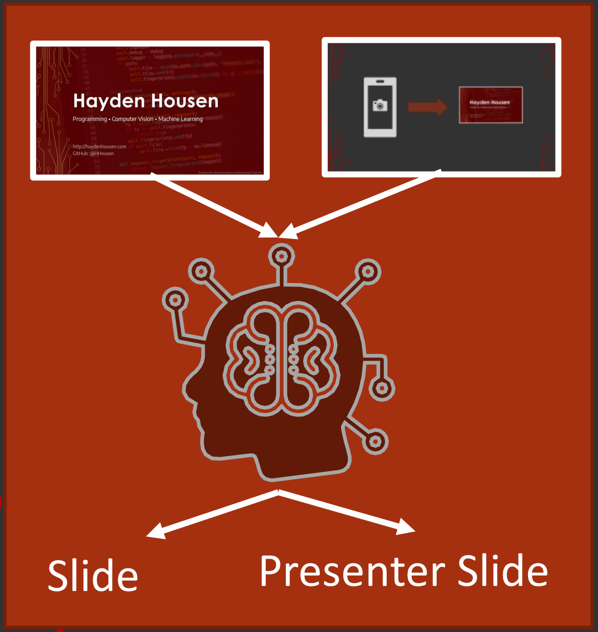

Slide Classifier

AI Summarization Models

End-to-End Process

Website

Main Components

Lecture Video Dataset

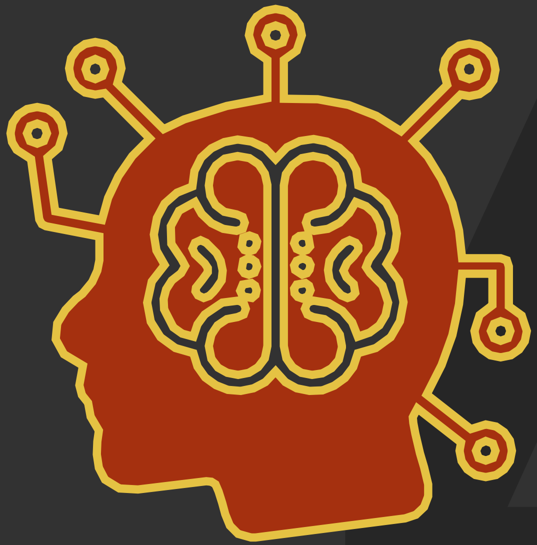

AI Slide Classifier Model

AI Summarization Models

End-to-End Process

Website

1. Lecture Video Dataset – Literature Review

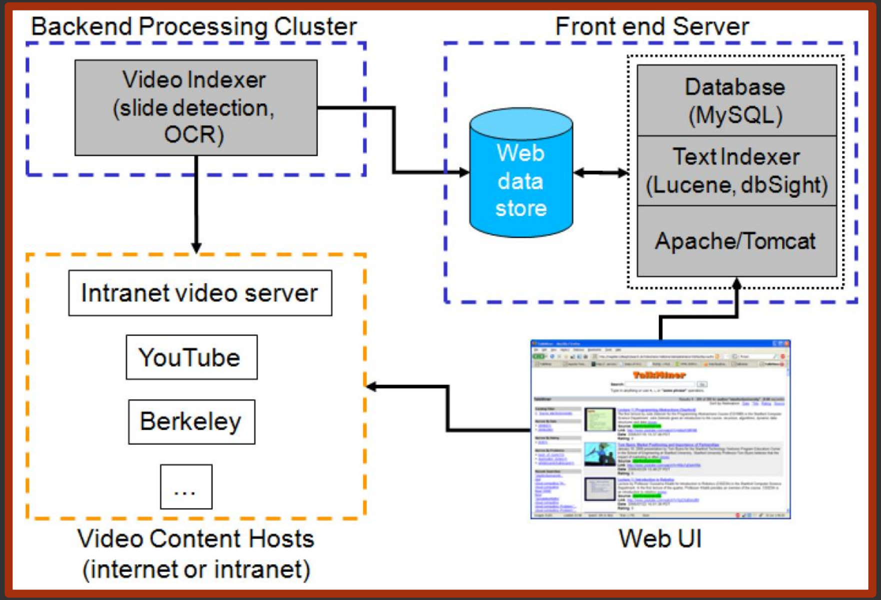

Summary of TalkMiner [22]

"TalkMiner: A lecture webcast search engine" by Adcock et al.[22] – Web scraping and OCR used to index online lecture videos

1. Lecture Video Dataset Diagram

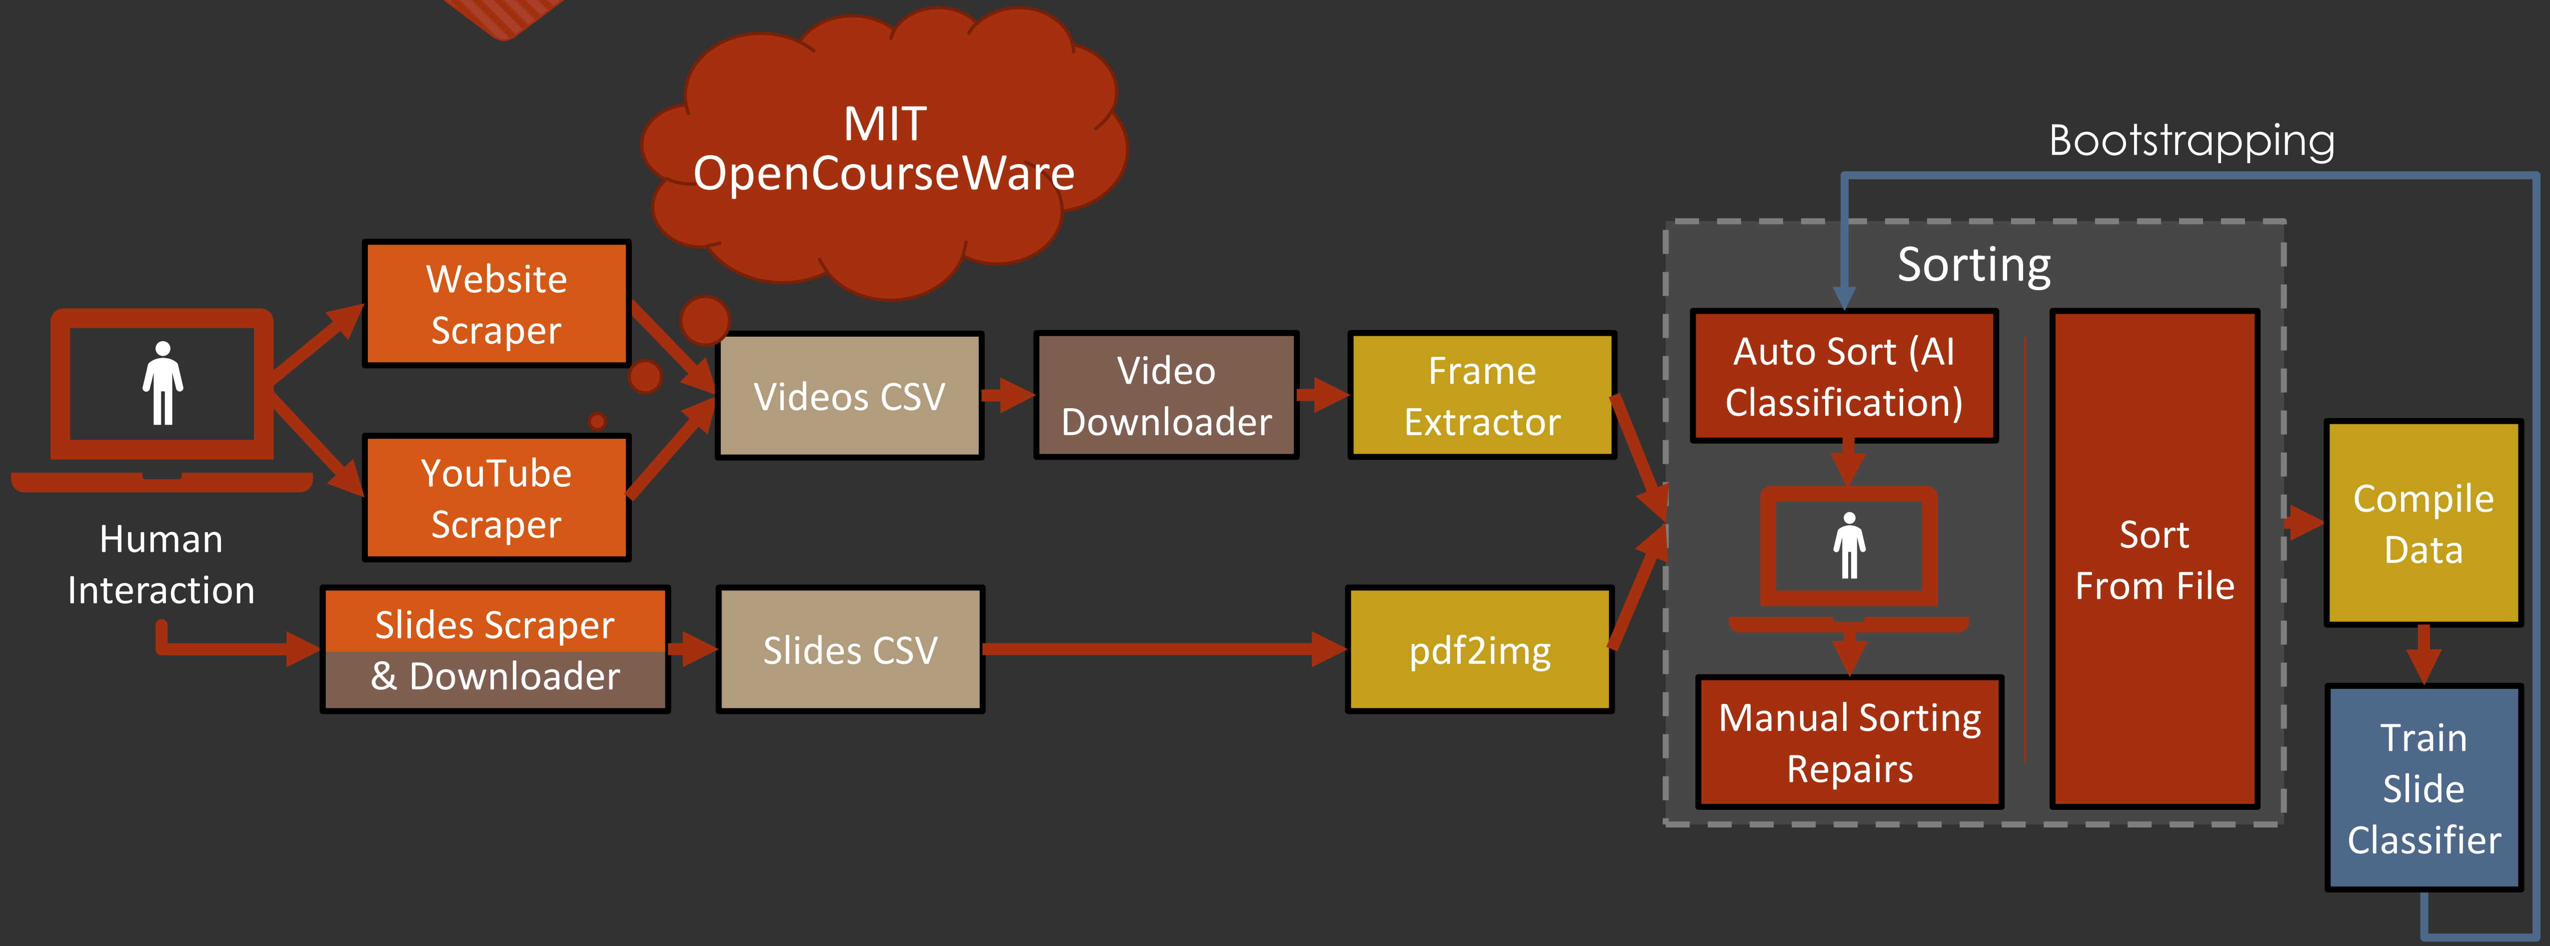

1. Lecture Video Dataset – Mass Collection

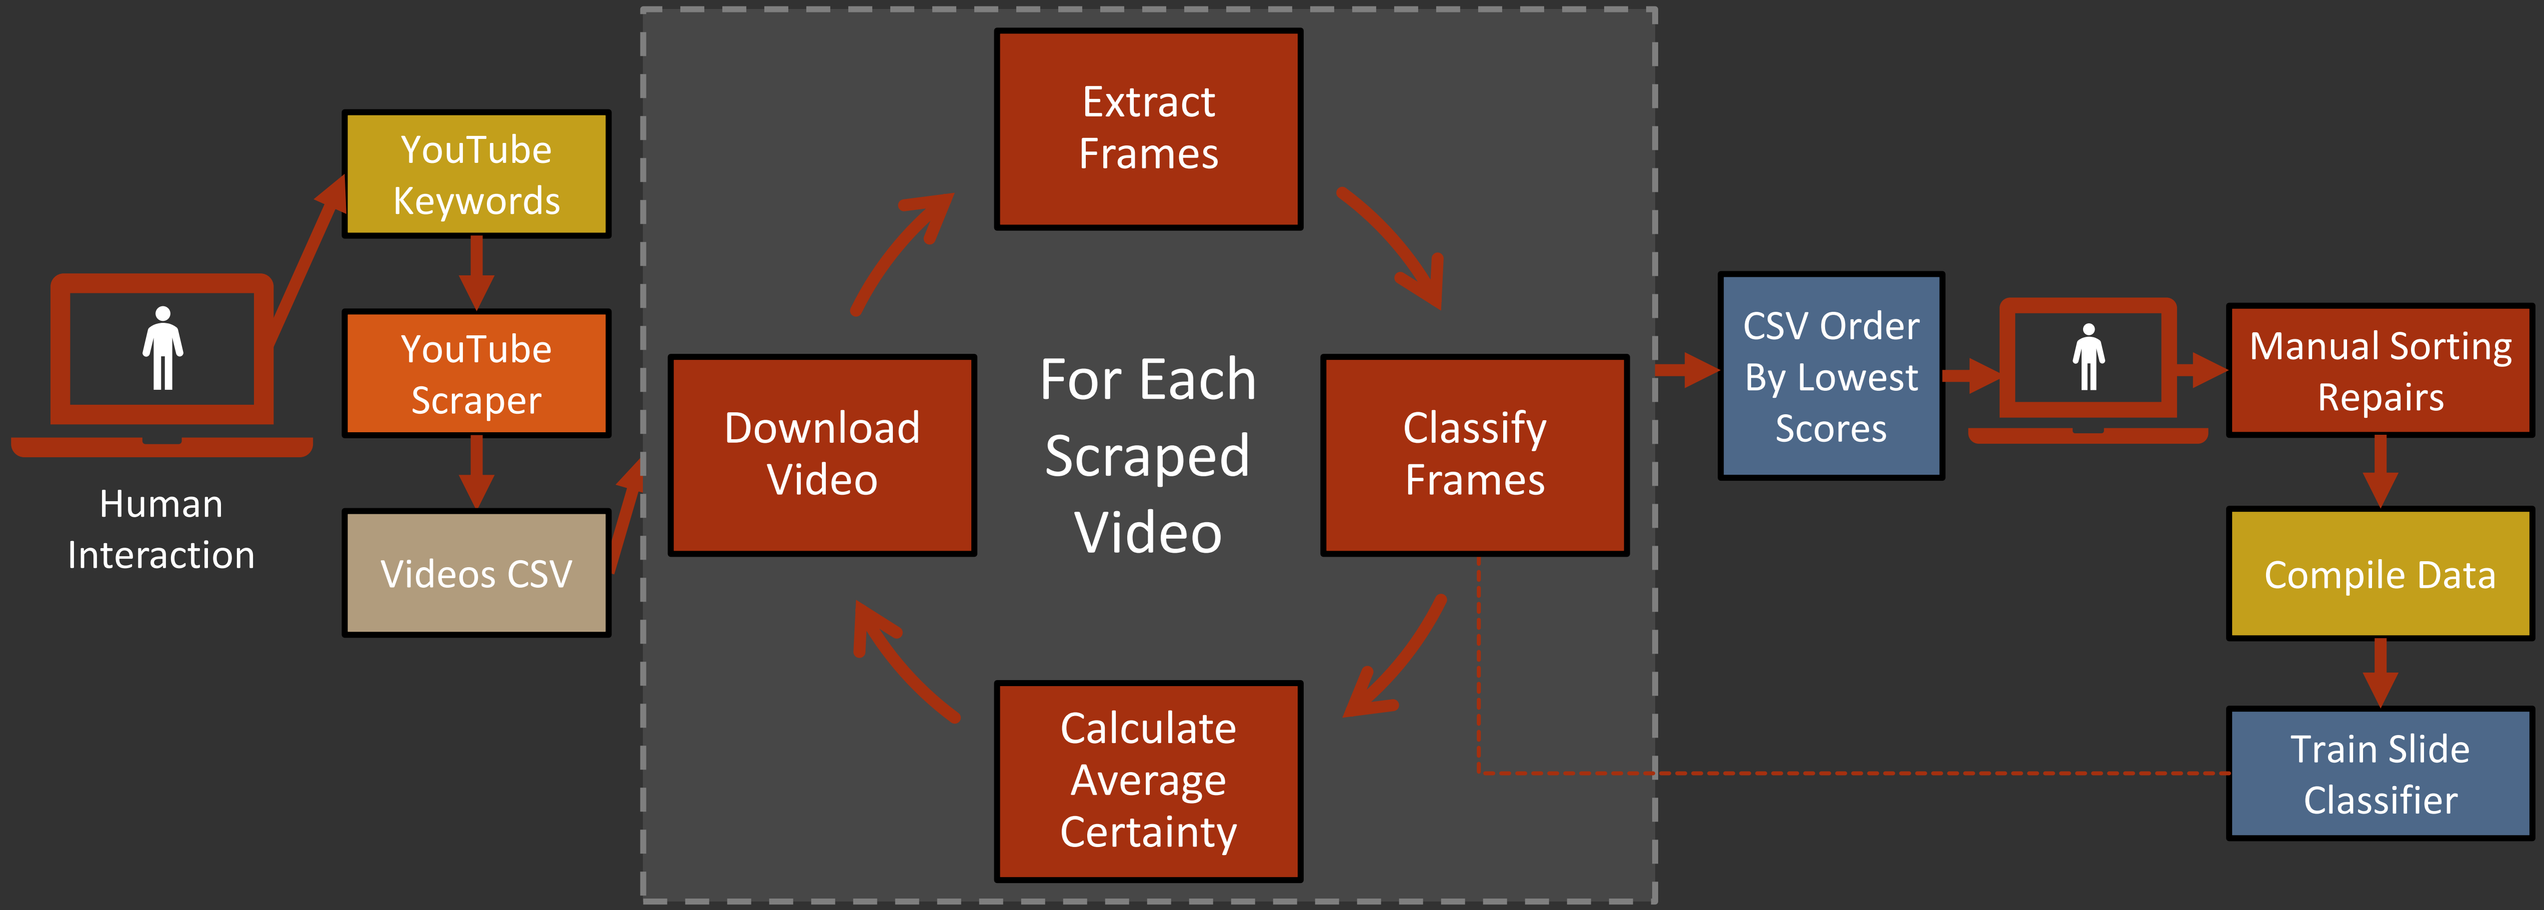

2. Slide Classifier – Literature Review

"Unsupervised Video Summarization with Adversarial LSTM Networks"

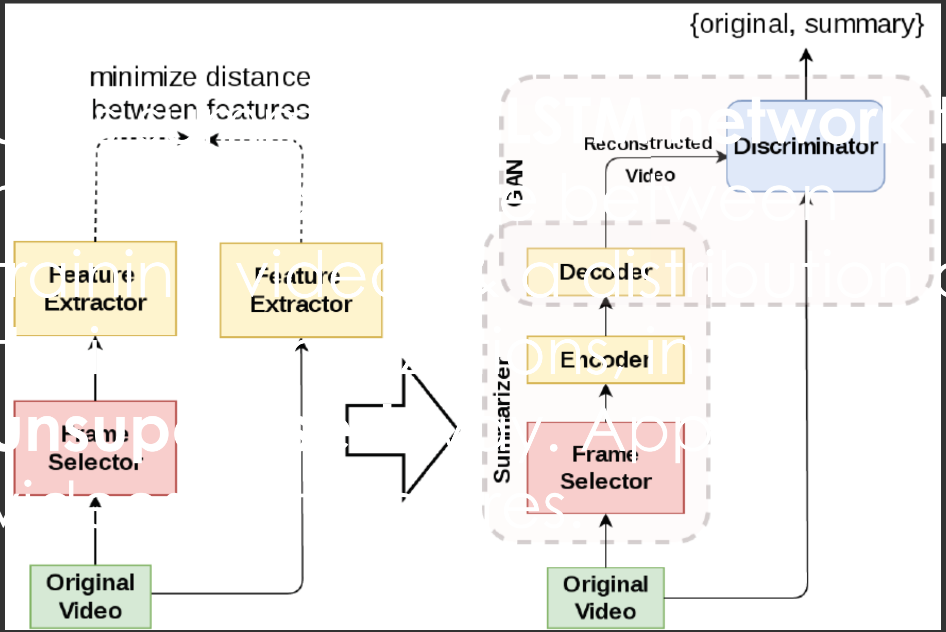

Summarizer network from [19]

"Lecture Video Indexing and Analysis Using Video OCR Technology"

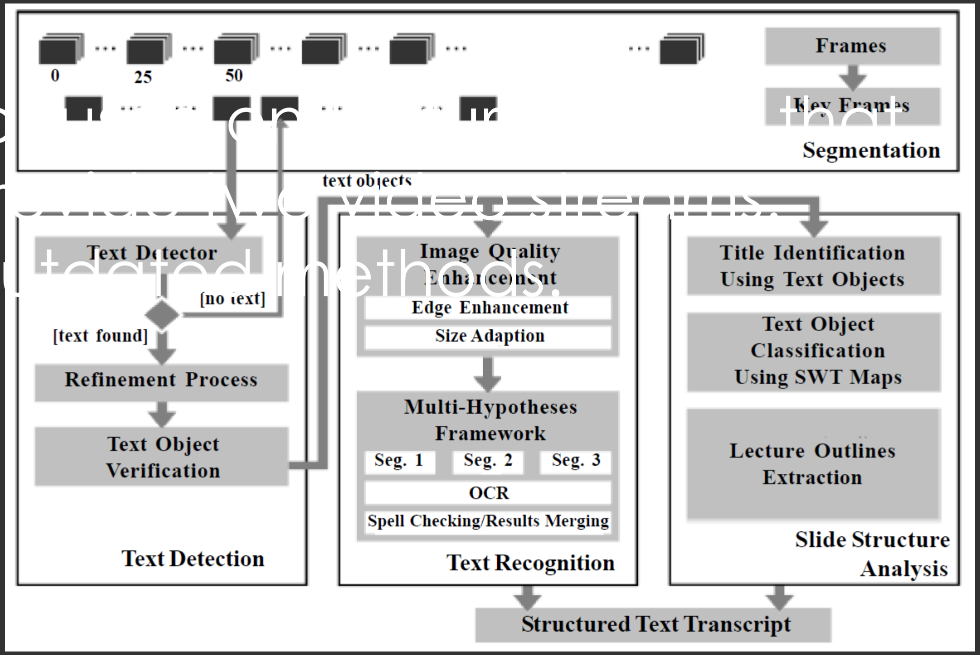

Entire System Workflow of [20]

2. Slide Classifier – ML Workflow

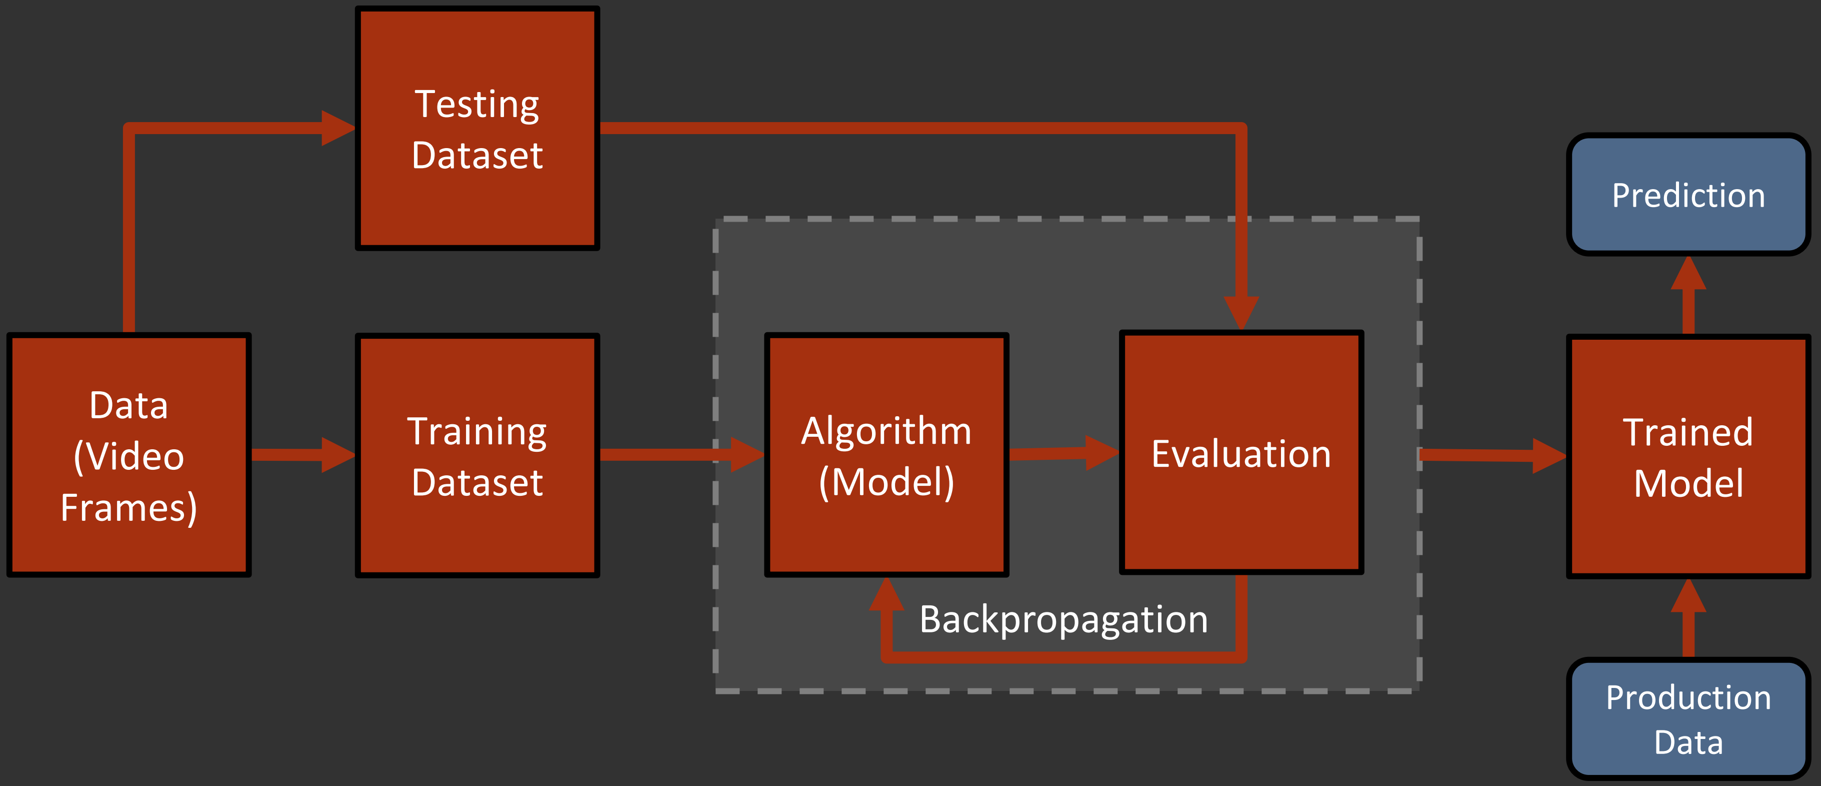

Classes
Audience
Audience_presenter
Audience_presenter_slide
Demo
Presenter
Presenter_slide
Presenter_whiteboard
Slide
Whiteboard

2. Slide Classifier – Training Code

slide-classifier-pytorch.py

```
import argparse
import os
import random
import shutil
import time
import warnings

import torch
import torch.nn as nn
import torch.nn.parallel
import torch.backends.cudnn as cudnn
import torch.distributed as dist
import torch.optim
import torch.multiprocessing as mp
import torch.utils.data
import torch.utils.data.distributed
import torchvision.transforms as transforms
import torchvision.datasets as datasets
import torchvision.models as models

import matplotlib as mpl
```

2. Slide Classifier – Algorithm (ResNet-34)

[26]

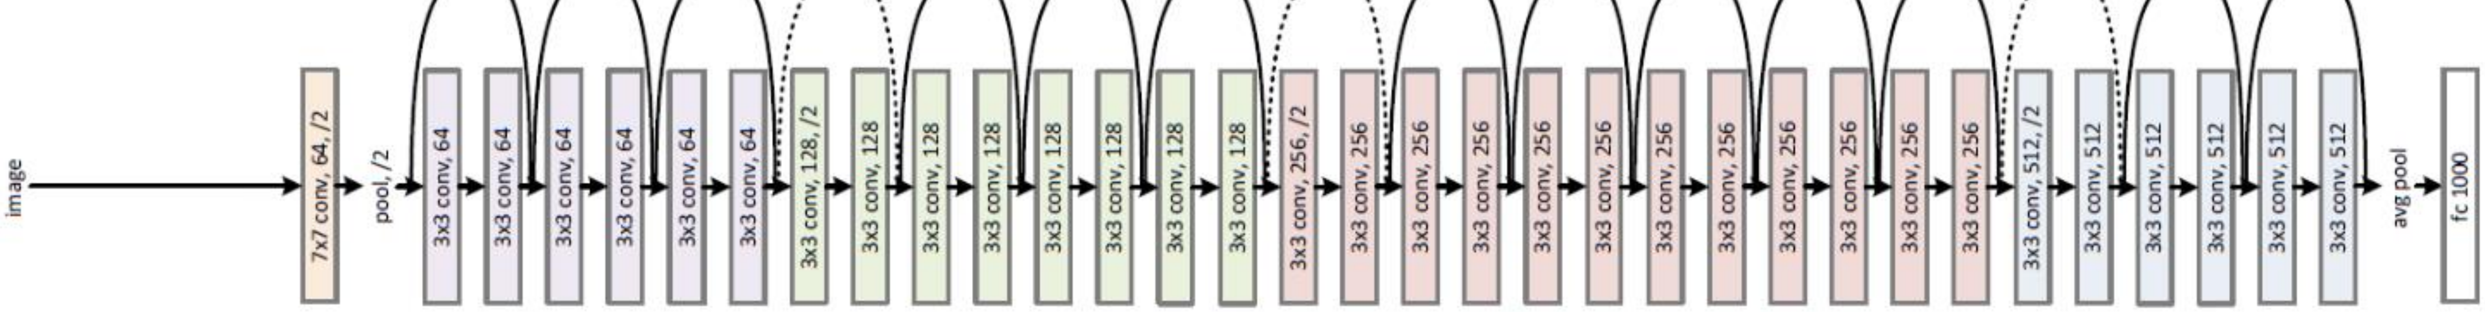

2. Slide Classifier – Algorithm (ResNet-34)

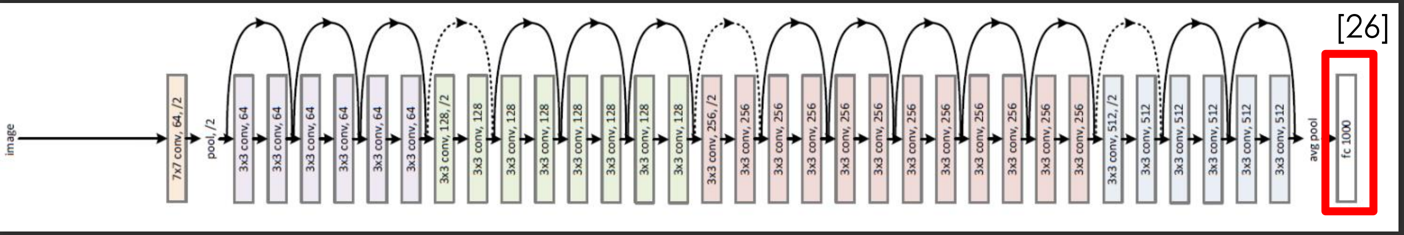
[26]

- **ResNet** is a type of Convolutional Neural Network (**CNN**), NN that interprets images
- **CNN Simple Definition:** Series of matrix multiplications
- I've added some regularization methods (improve accuracy)

```
layers = [
    AdaptiveConcatPool2d(1),
    Flatten(),
    nn.BatchNorm1d(1024),
    nn.Dropout(0.25),
    nn.Linear(1024, 512),
    nn.ReLU(inplace=True),
    nn.BatchNorm1d(512),
    nn.Dropout(0.5),
    nn.Linear(512, num_out_features)
]
```

Convolutional Neural Networks

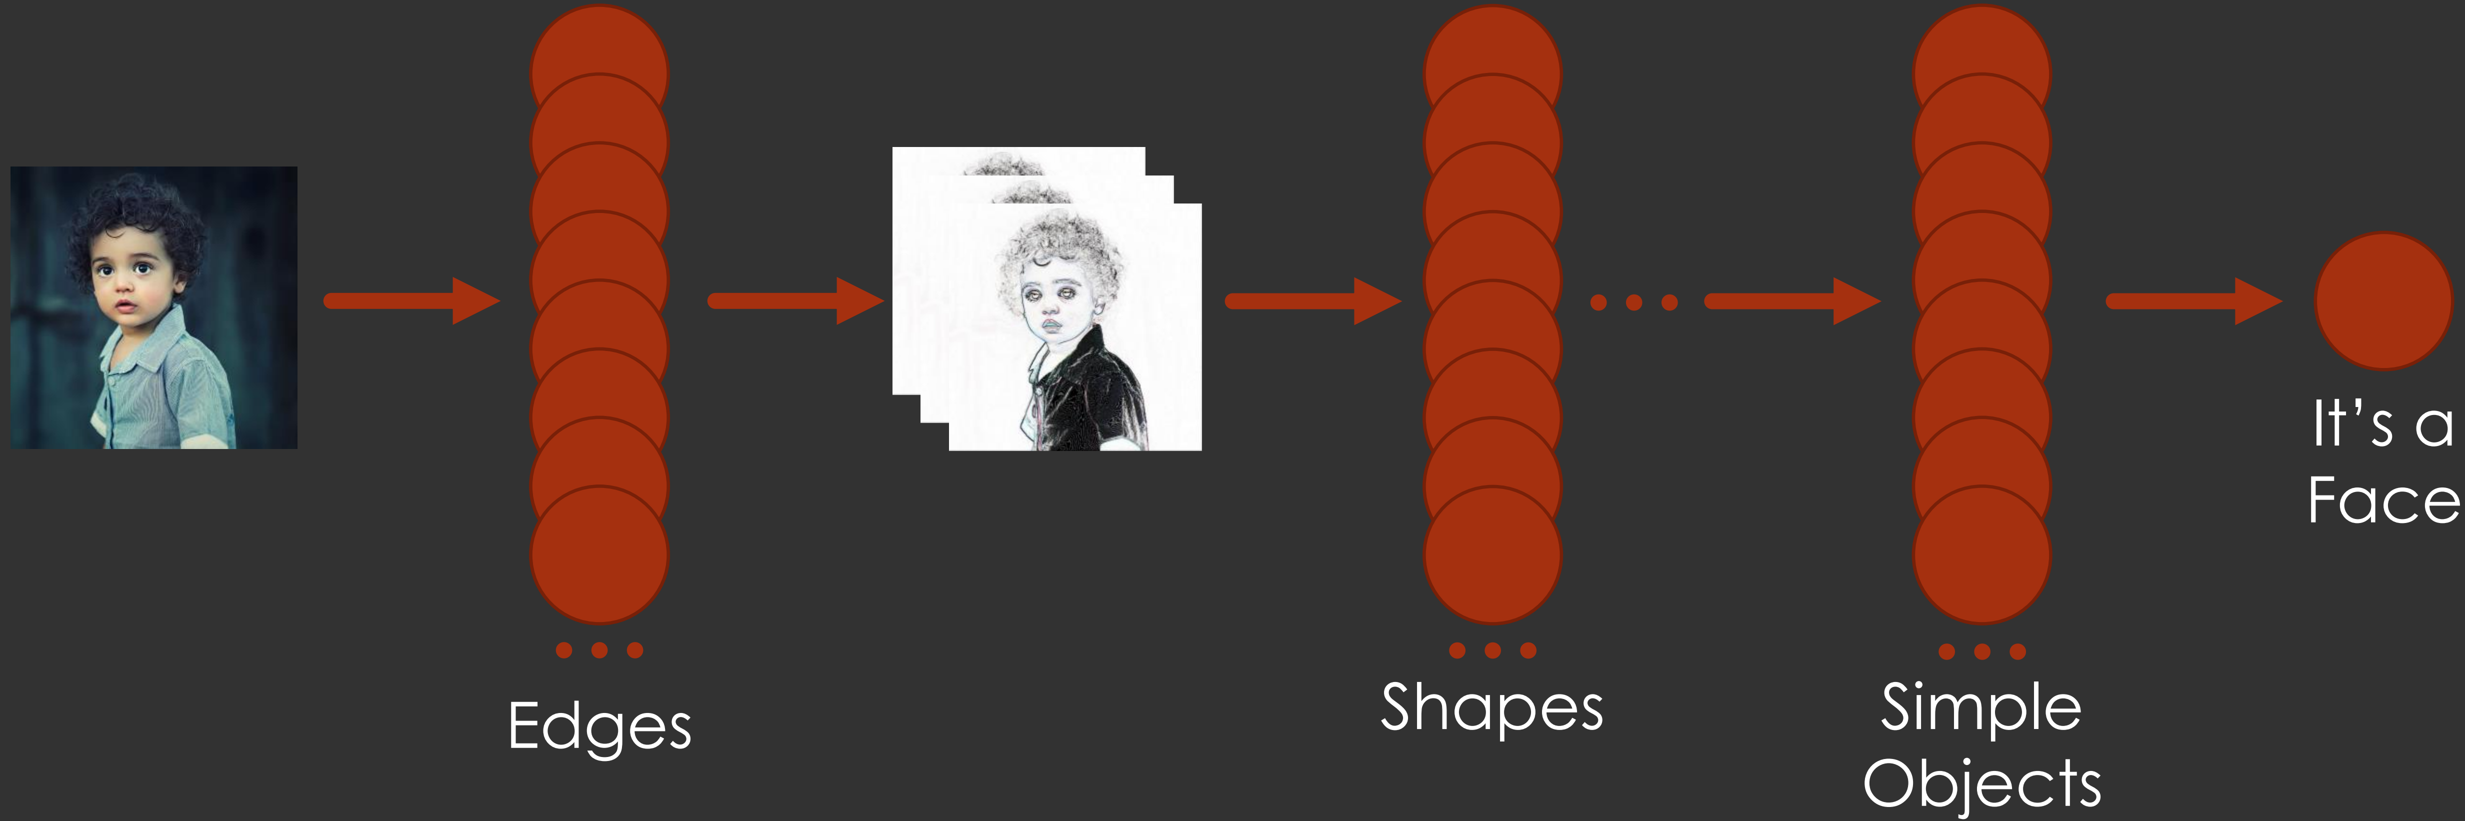

https://pixabay.com/en/child-boy-portrait-children-people-817373/

Slide Classifier – Current Results

	precision	recall	f1-score	support
audience	0.00	0.00	0.00	2
audience_presenter	0.84	0.90	0.87	30
audience_presenter_slide	1.00	0.43	0.61	23
demo	0.98	0.93	0.95	45
presenter	0.98	0.99	0.99	582
presenter_slide	0.97	1.00	0.99	626
presenter_whiteboard	0.95	0.95	0.95	60
slide	1.00	0.99	1.00	442
whiteboard	0.96	0.88	0.92	25
accuracy			0.98	1835
macro avg	0.85	0.79	0.81	1835
weighted avg	0.98	0.98	0.98	1835

Slide Classifier – Current Results

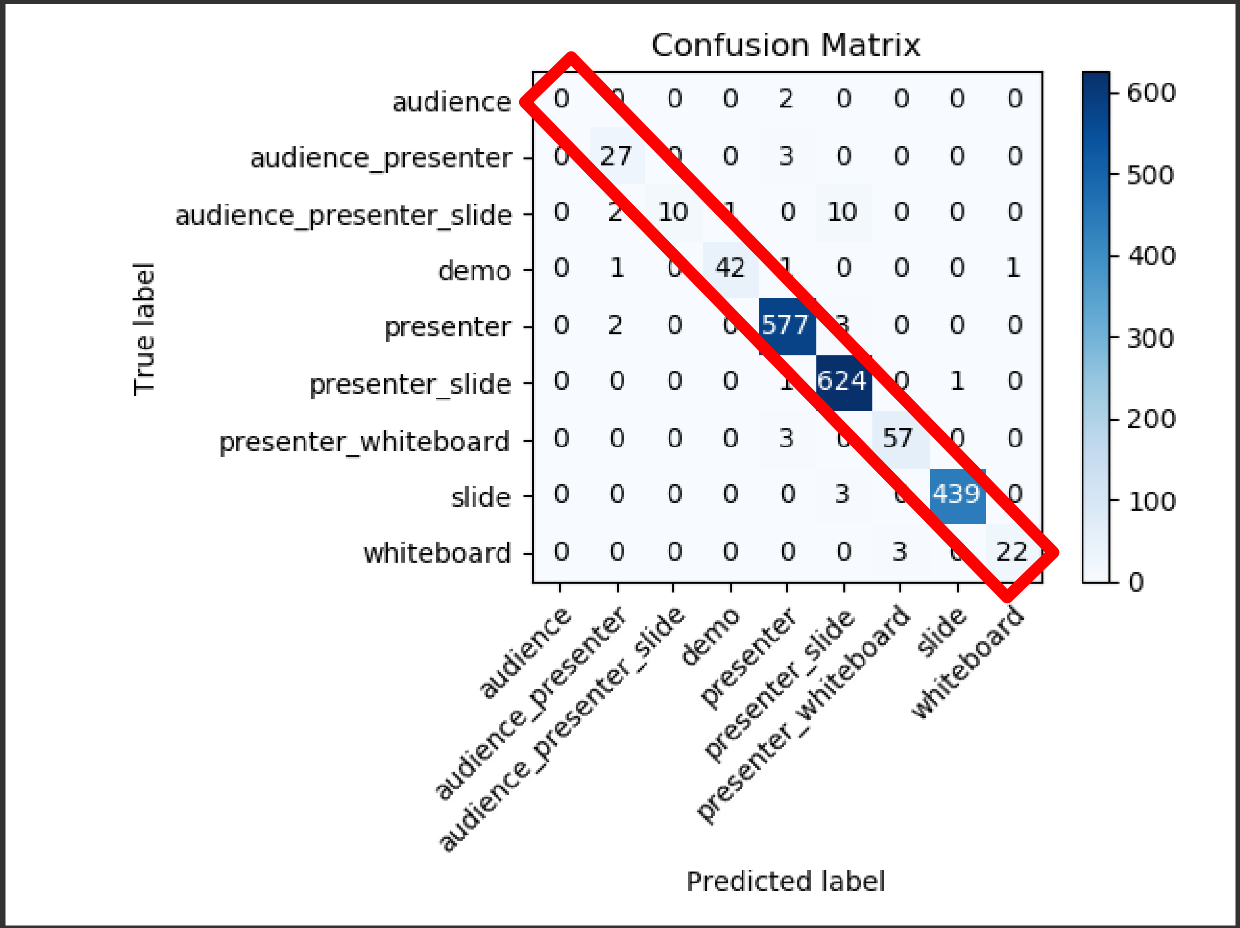

3. End-to-End Approach – Overview

Extract Frames → **Classify Frames** → **Perspective Crop** → **Cluster Slides** →

Extract Frames	Classify Frames	Perspective Crop	Cluster Slides
ffmpeg	Slide Classifier that uses the ML model from previous slide trained on the dataset	Crop slides captured by video camera	Segment or Normal

3. End-to-End Approach – Overview

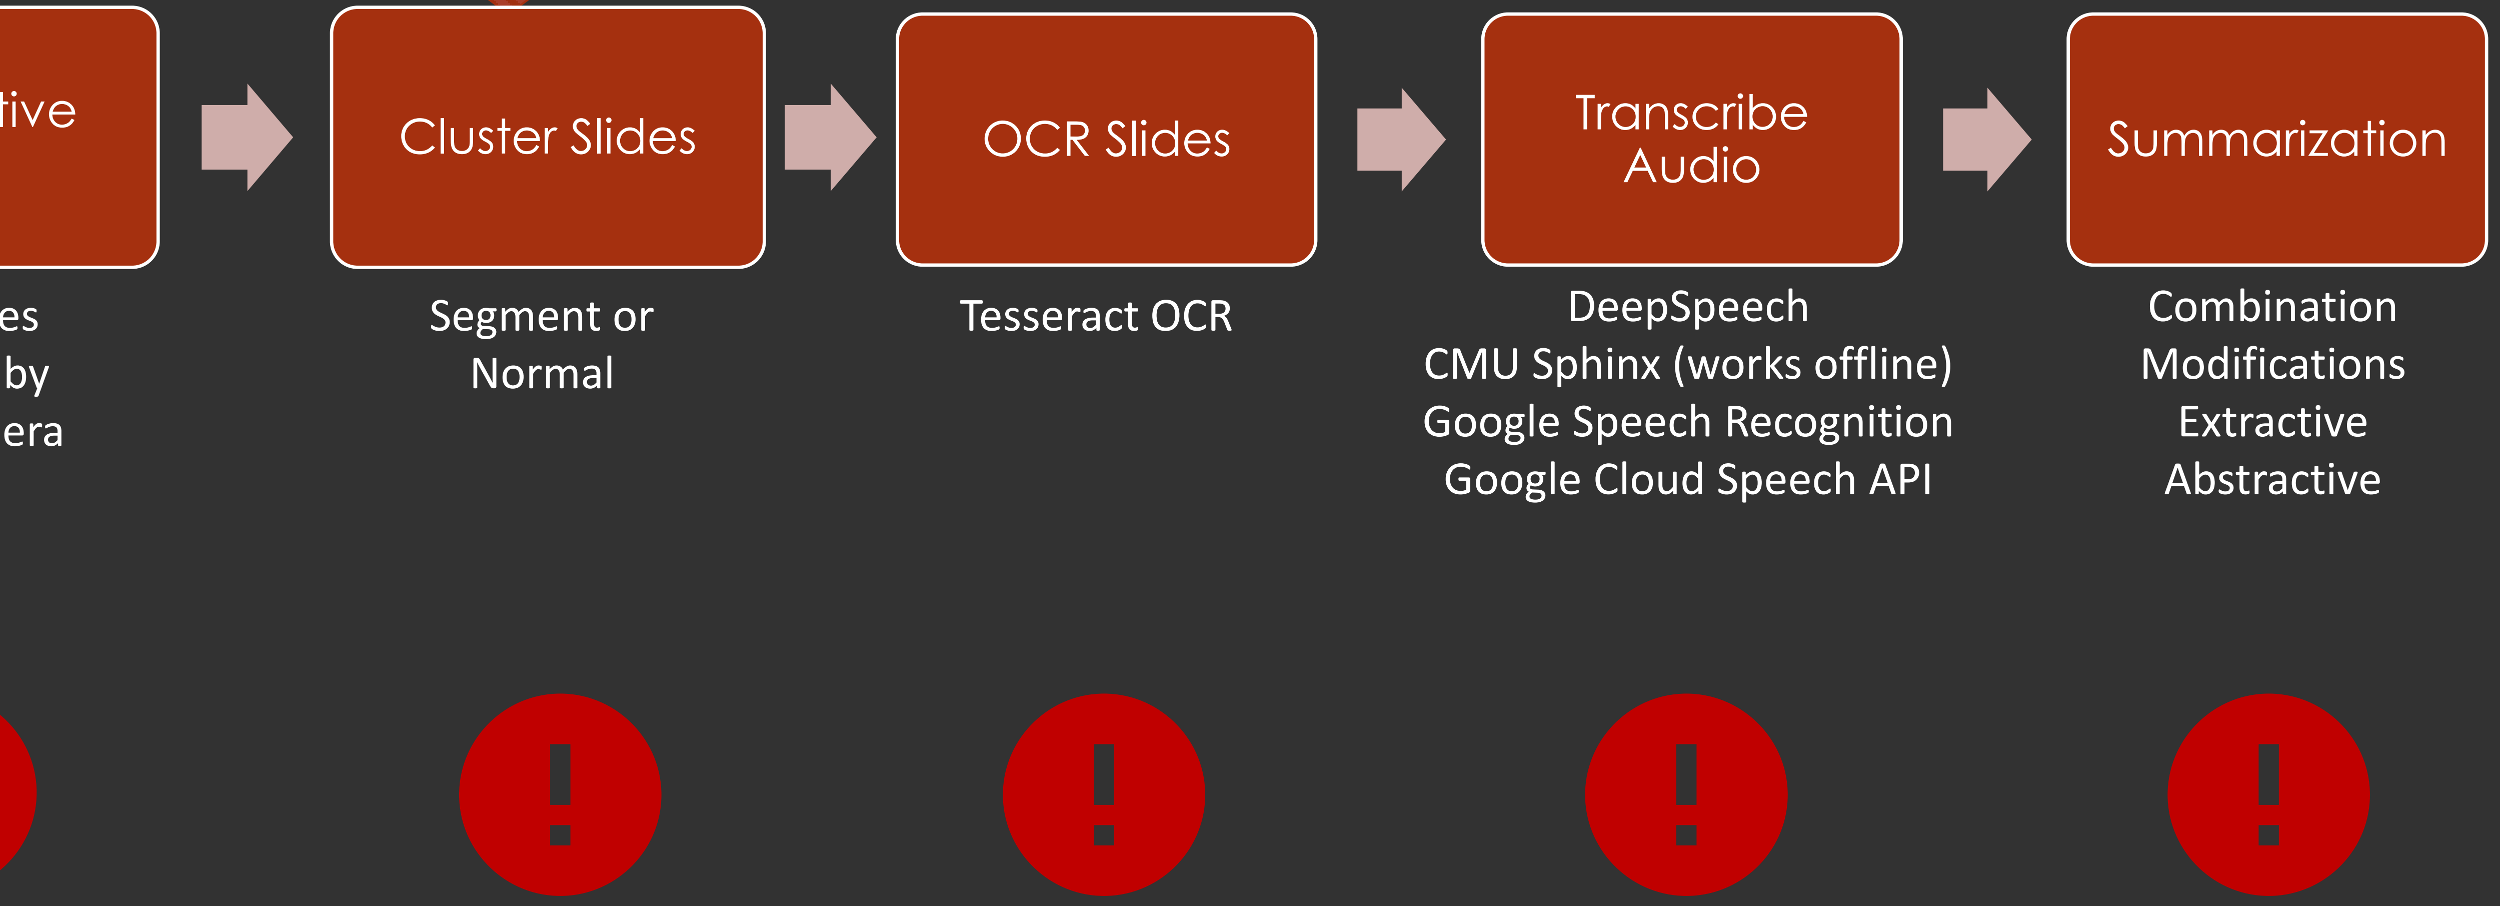

Perspective Cropping

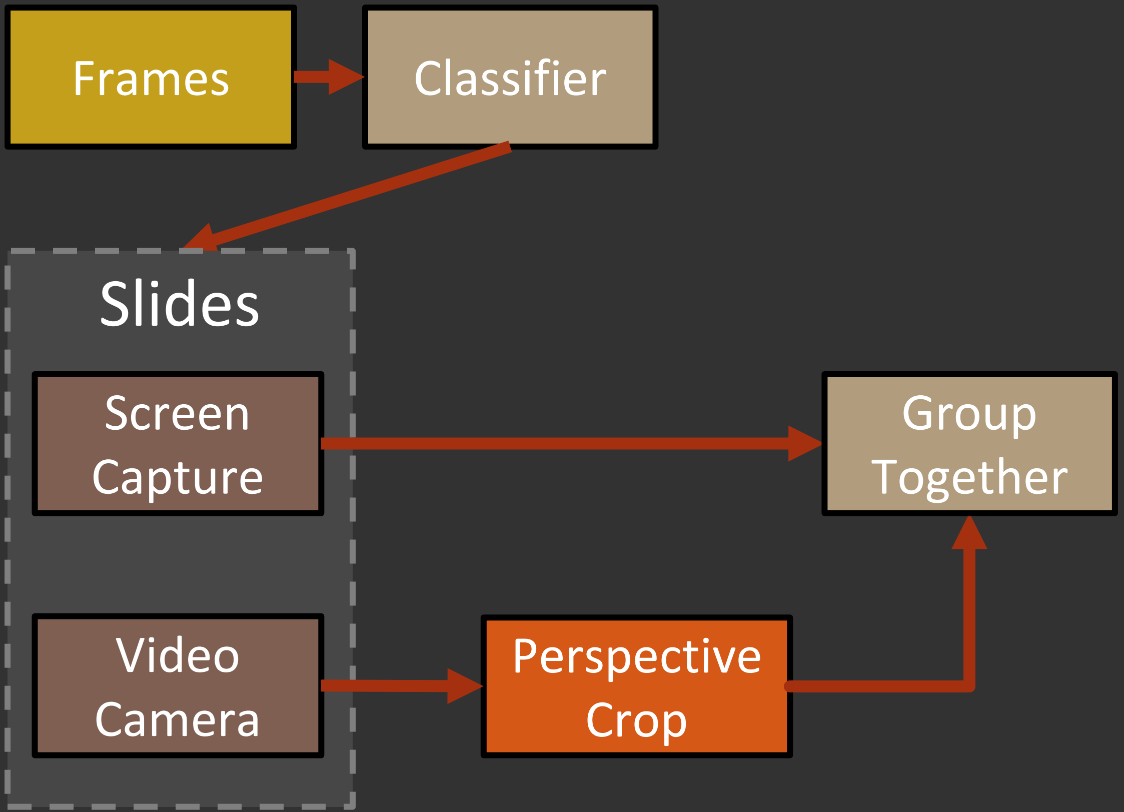

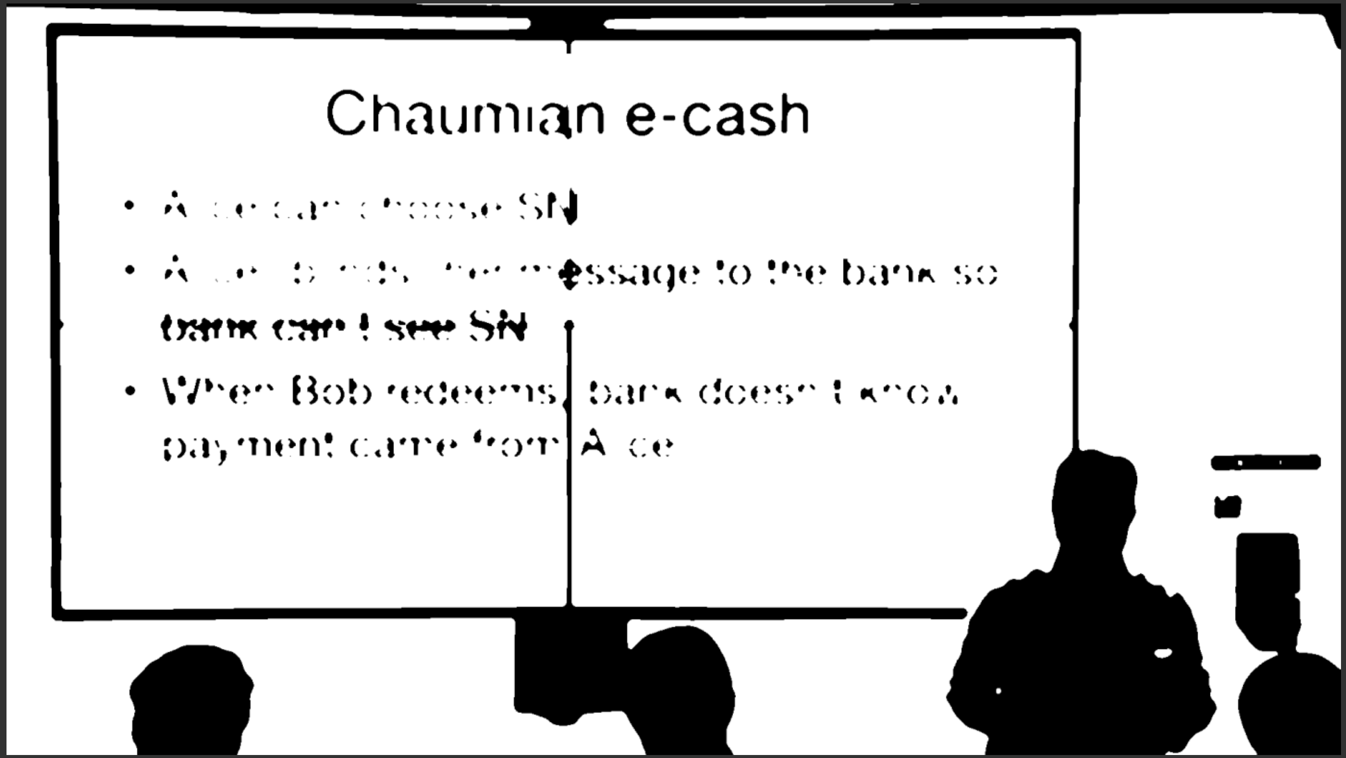

Perspective Cropping

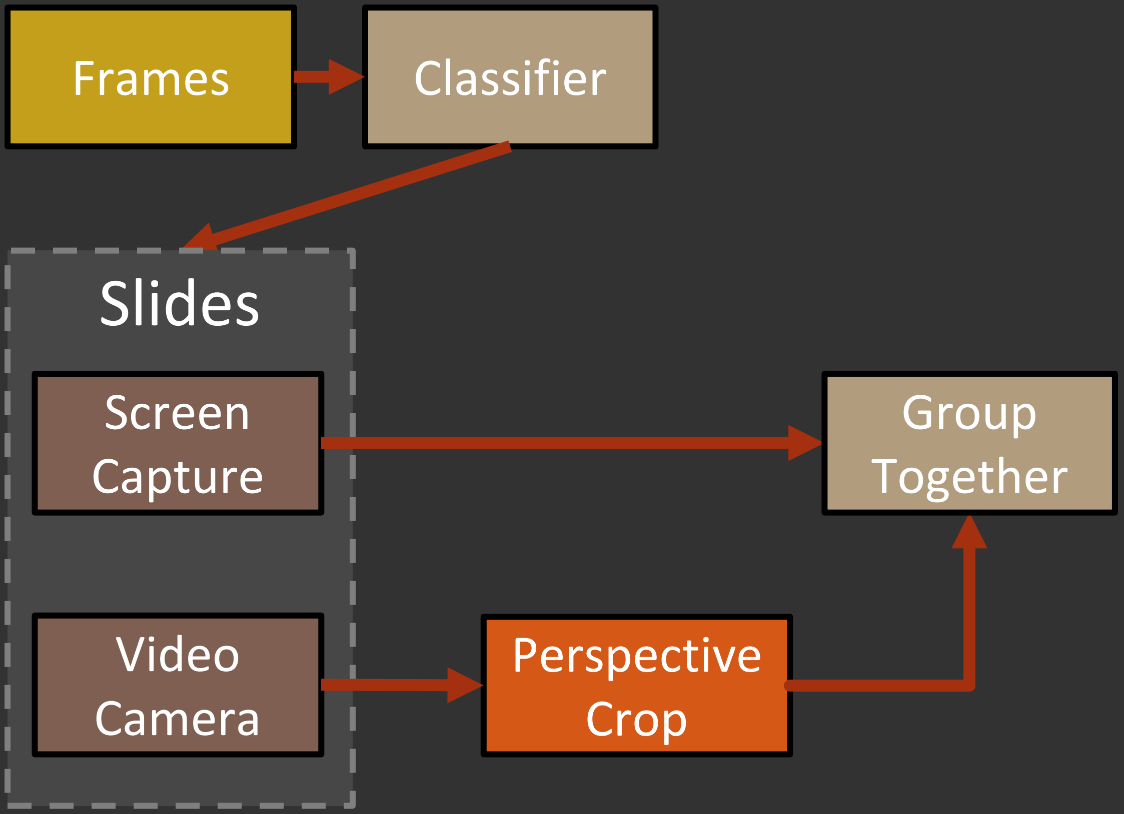

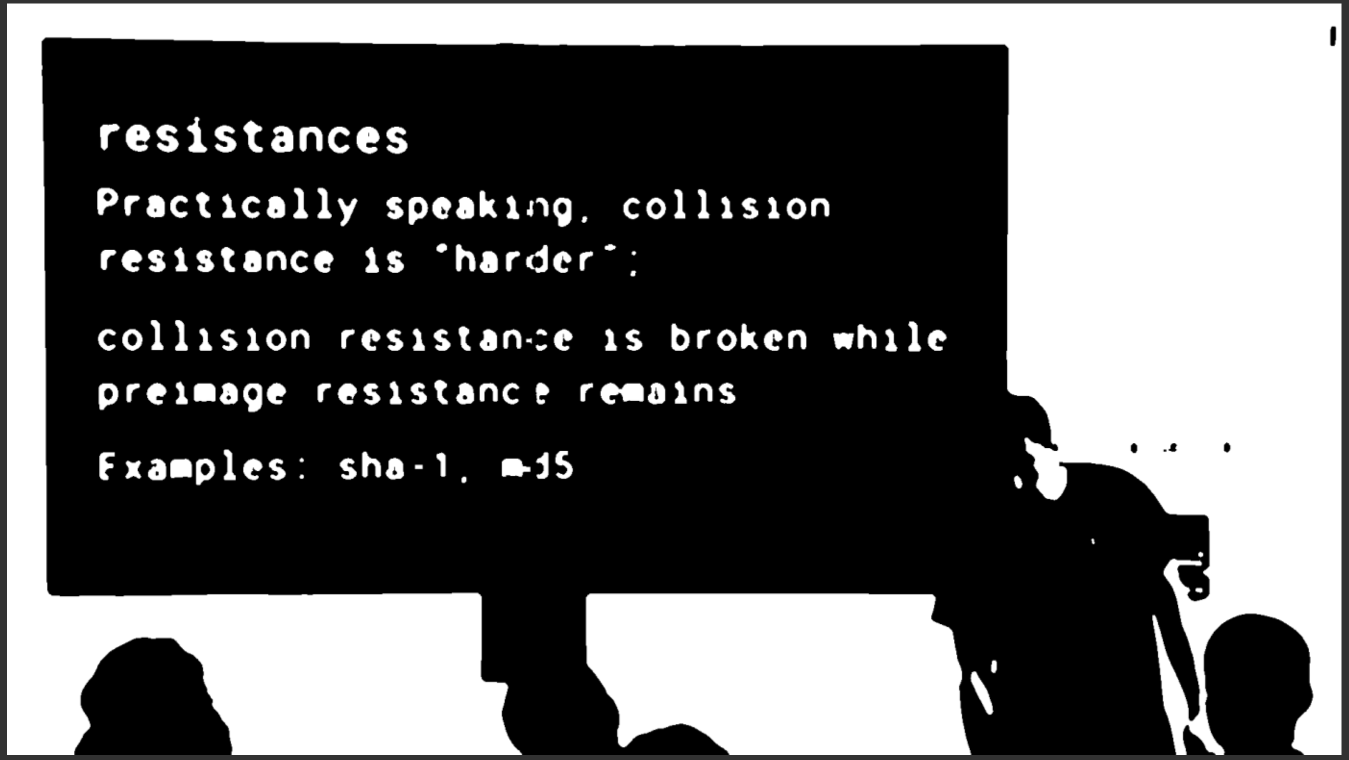

Perspective Cropping – Problems

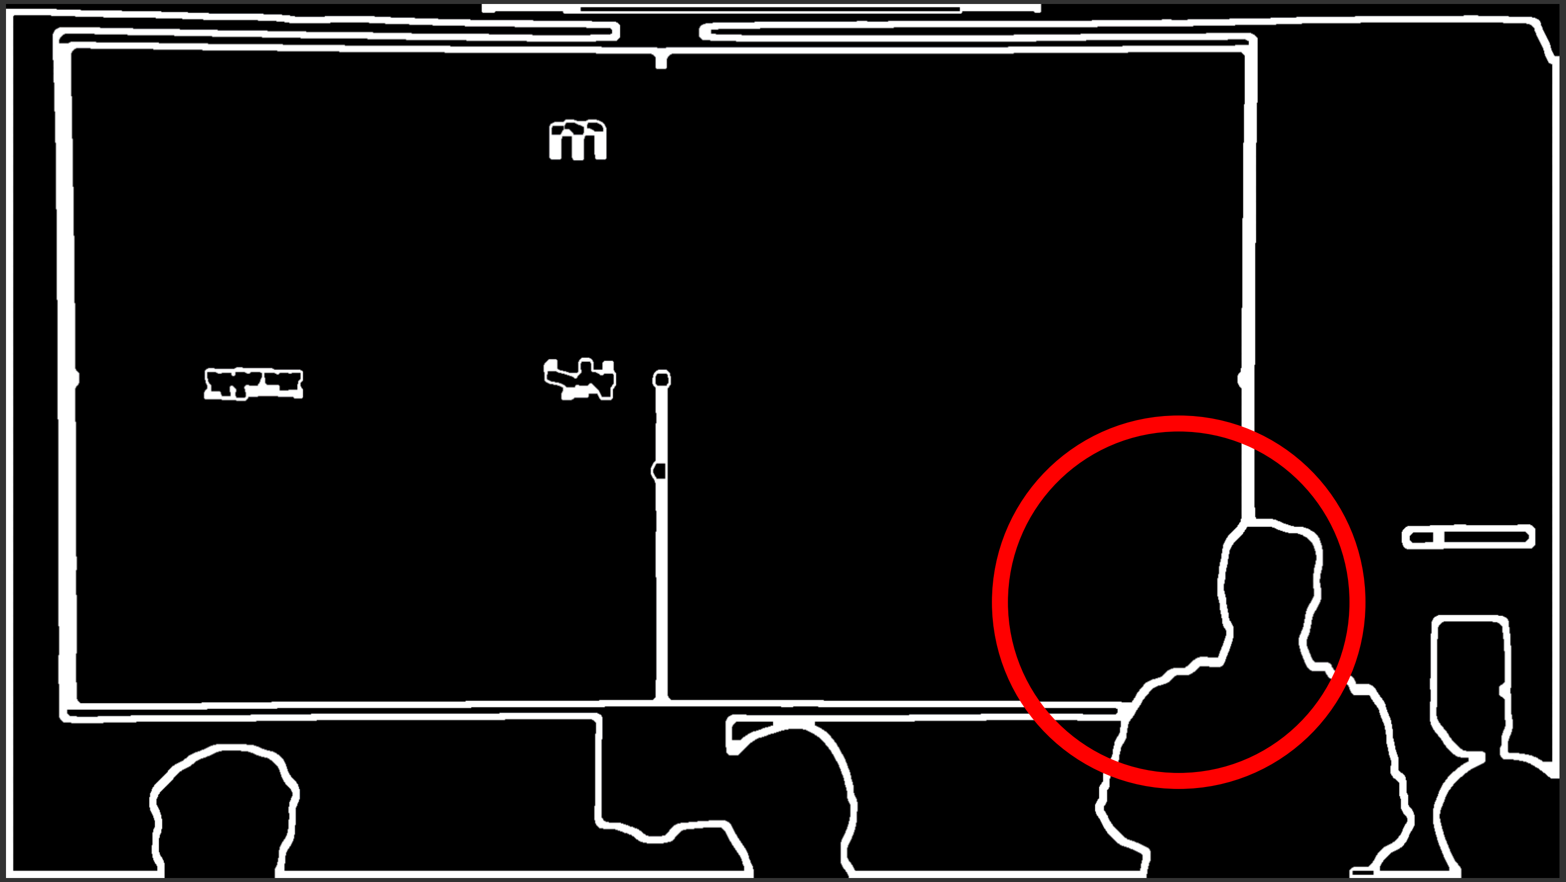

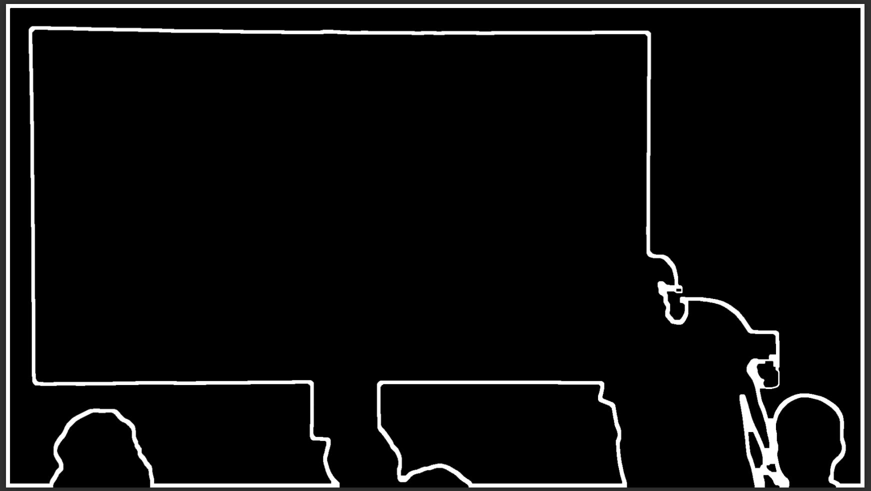

Perspective Cropping – Problems

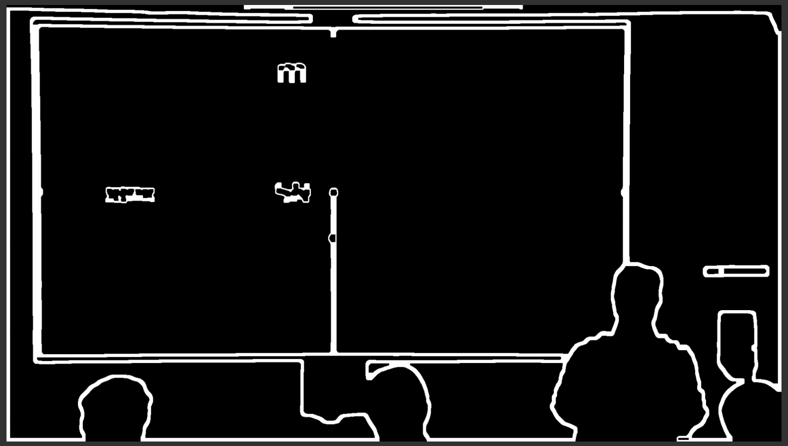

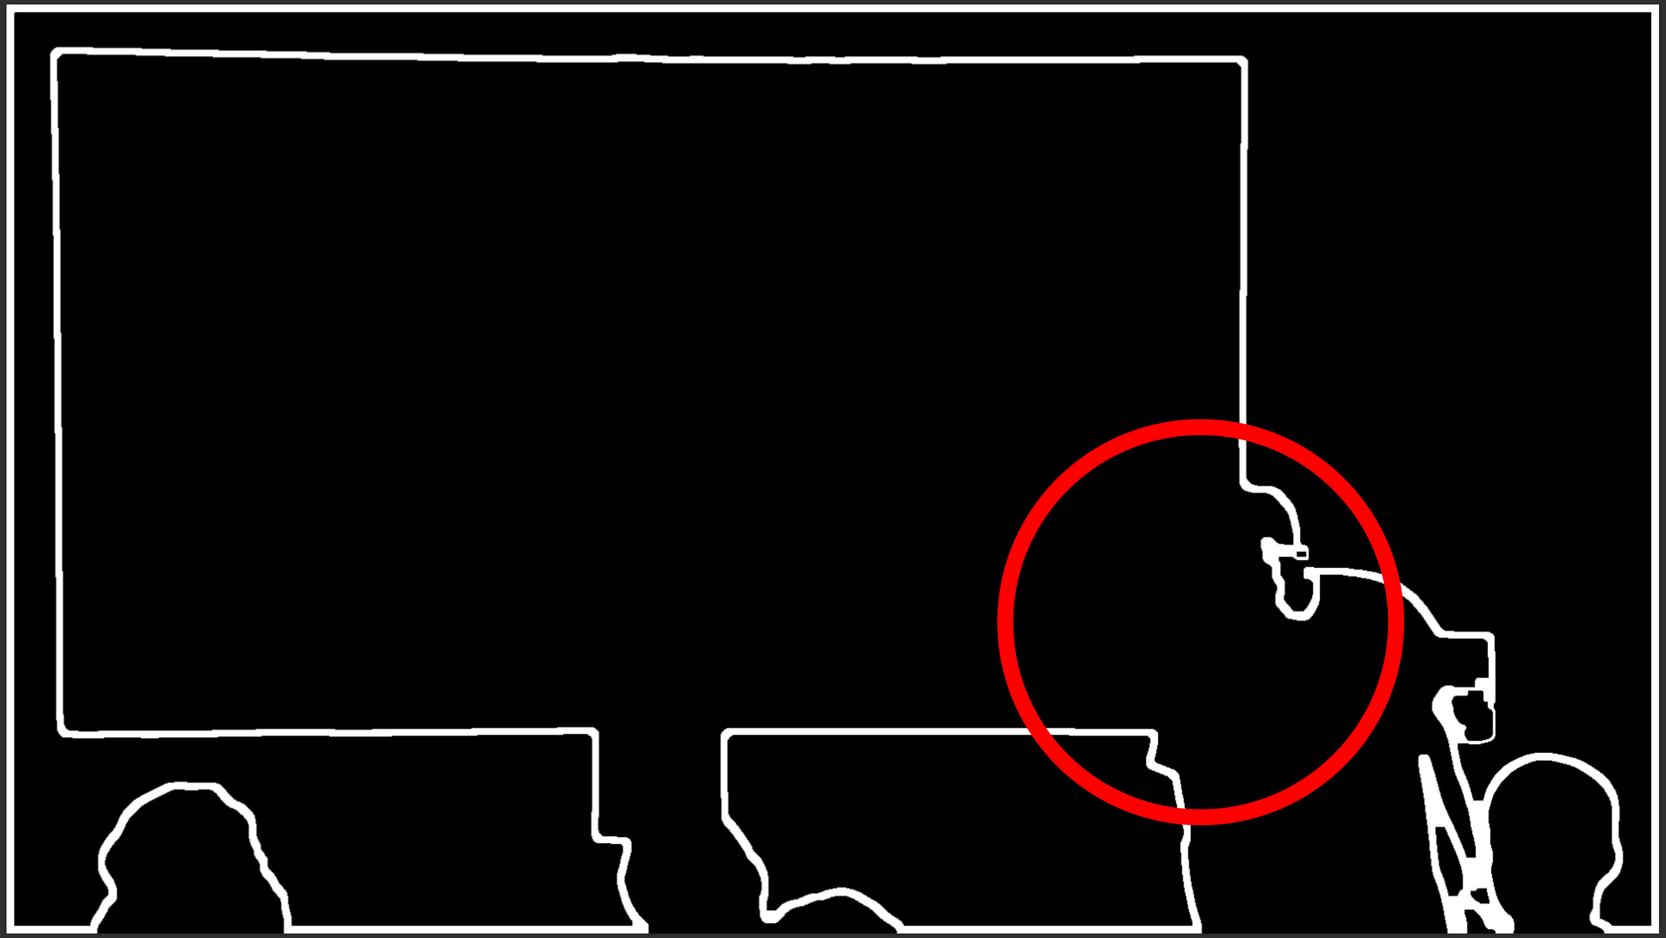

Clustering

Normal

- **Groups slides** based on **features** (visual similarities) extracted from *Slide Classifier CNN*
- *Affinity Propagation* and *K-Means*
- Similar slides are clustered together (removes transitions)
- Eliminates **duplicate frames** that show the **same slide**

cluster.py

```python
import sys, os, shutil
from tqdm import tqdm
from helpers import make_dir_if_not_exist
# Hack to import modules from different parent directory
sys.path.insert(1, os.path.join(sys.path[0], '../Models/slide-classifier'))
from class_cluster_scikit import Cluster
from custom_nnmodules import *
from inference import *

def make_clusters(slides_dir, copy=True):
    """Clusters all images in directory `slides_dir` and saves each cluster to a subfolder in
`cluster_dir` (directory in parent of `slides_dir`)"""
    slides = os.listdir(slides_dir)
    num_slides = len(slides)
    cluster = Cluster(algorithm_name="affinity_propagation", preference=-8, damping=0.72)

    print("> AI Clustering Engine: Ready to cluster " + str(num_slides) + " slides")
    for idx, slide in tqdm(enumerate(slides), total=num_slides, desc="> AI Clustering Engine: Feature
extraction"):
        current_slide_path = os.path.join(slides_dir, slide)
        _, _, _, extracted_features = get_prediction(Image.open(current_slide_path))
        cluster.add(extracted_features, slide)

    #cluster.calculate_best_k()
    move_list = cluster.create_move_list()

    num_clusters = cluster.get_num_clusters()
    print("Predicted Number of Clusters: " + str(num_clusters))

    cluster_dir = slides_dir.parents[0] / "slide_clusters" # cluster_dir = up one directory from
slides_dir then into "slide_clusters"
    for filename in tqdm(move_list, desc="> AI Clustering Engine: Move/copy into cluster folders"):
        cluster_number = move_list[filename]
        current_slide_path = os.path.join(slides_dir, filename)
        current_cluster_path = cluster_dir / str(cluster_number)
        make_dir_if_not_exist(current_cluster_path)
        if copy:
            shutil.copy(str(current_slide_path), str(current_cluster_path))
        else:
            shutil.move(str(current_slide_path), str(current_cluster_path))

    return cluster_dir
```

Clustering

Normal

- **Groups slides** based on **features** (visual similarities) extracted from *Slide Classifier CNN*
- *Affinity Propagation* and *K-Means*
- Similar slides are clustered together (removes transitions)
- Eliminates **duplicate frames** that show the **same slide**

Segment

- **Iterates** through the extracted **slides in order**
- Marks a split when the **cosine similarity** between the feature vectors differs by a **value greater than the mean** of the cosine similarities

OCR

- *pytesseract* Python Package
- "**Optical character recognition** (OCR) tool for python"
- **Recognizes** and "**reads**" the **text** embedded in **images**
- Uses **Google's** Tesseract-OCR Engine

ocr.py

```
def all_in_folder(path):
    """Perform OCR on every file in folder and return results"""
    results = []
    images = os.listdir(path)
    images.sort()
    for item in tqdm(images, total=len(images), desc="> OCR: Progress"):
        print("> OCR: Processing file " + item)
        current_path = os.path.join(path, item)
        if os.path.isfile(current_path):
            ocr_result = pytesseract.image_to_string(Image.open(current_path))
            results.append(ocr_result)
    print("> OCR: Returning results")
    return results

def write_to_file(results, save_file):
    """Write everything stored in `results` to file at path `save_file`. Used to write results from
`all_in_folder()` to `save_file`."""
    file_results = open(save_file, "a+")
    print("> OCR: Writing results to file " + str(save_file))
    for item in tqdm(results, total=len(results), desc="> OCR: Writing To File Progress"):
        file_results.write(item + "\r\n")
    file_results.close()
    print("> OCR: Results written to " + str(save_file))
```

Transcribe Audio – YouTube

- If the **lecture** to be summarized is a **YouTube video**
- **Download** the transcript **directly** from YouTube
- Apply **minimal processing** (remove speaker names)
- Human-made transcripts improve summarization (less error from speech-to-text process)

Transcribe Audio – Speech-To-Text

- **DeepSpeech** architecture created by **Baidu** in **2014**
- **Project DeepSpeech** created by Mozilla (Firefox) to provide **open source** community with Speech-To-Text engine
- **5.97% word error rate** on the **LibriSpeech clean** test corpus (one of many speech datasets)
- Audio File (WAV) → Transcript (TXT)

https://mozilla.design/mozilla/logo-usage/

Transcribe Audio – Chunking

- **Voice Activity:** Uses WebRTC Voice Activity Detector (VAD) – reportedly one of the best VADs available, being fast, modern, free
- **Noise Activity:** Detects segments of audio that are significantly below the average loudness of the file.
- Chunking increases speed of voice-to-text by reducing the amount of audio without speech

3. End-to-End Approach – Slides/Videos

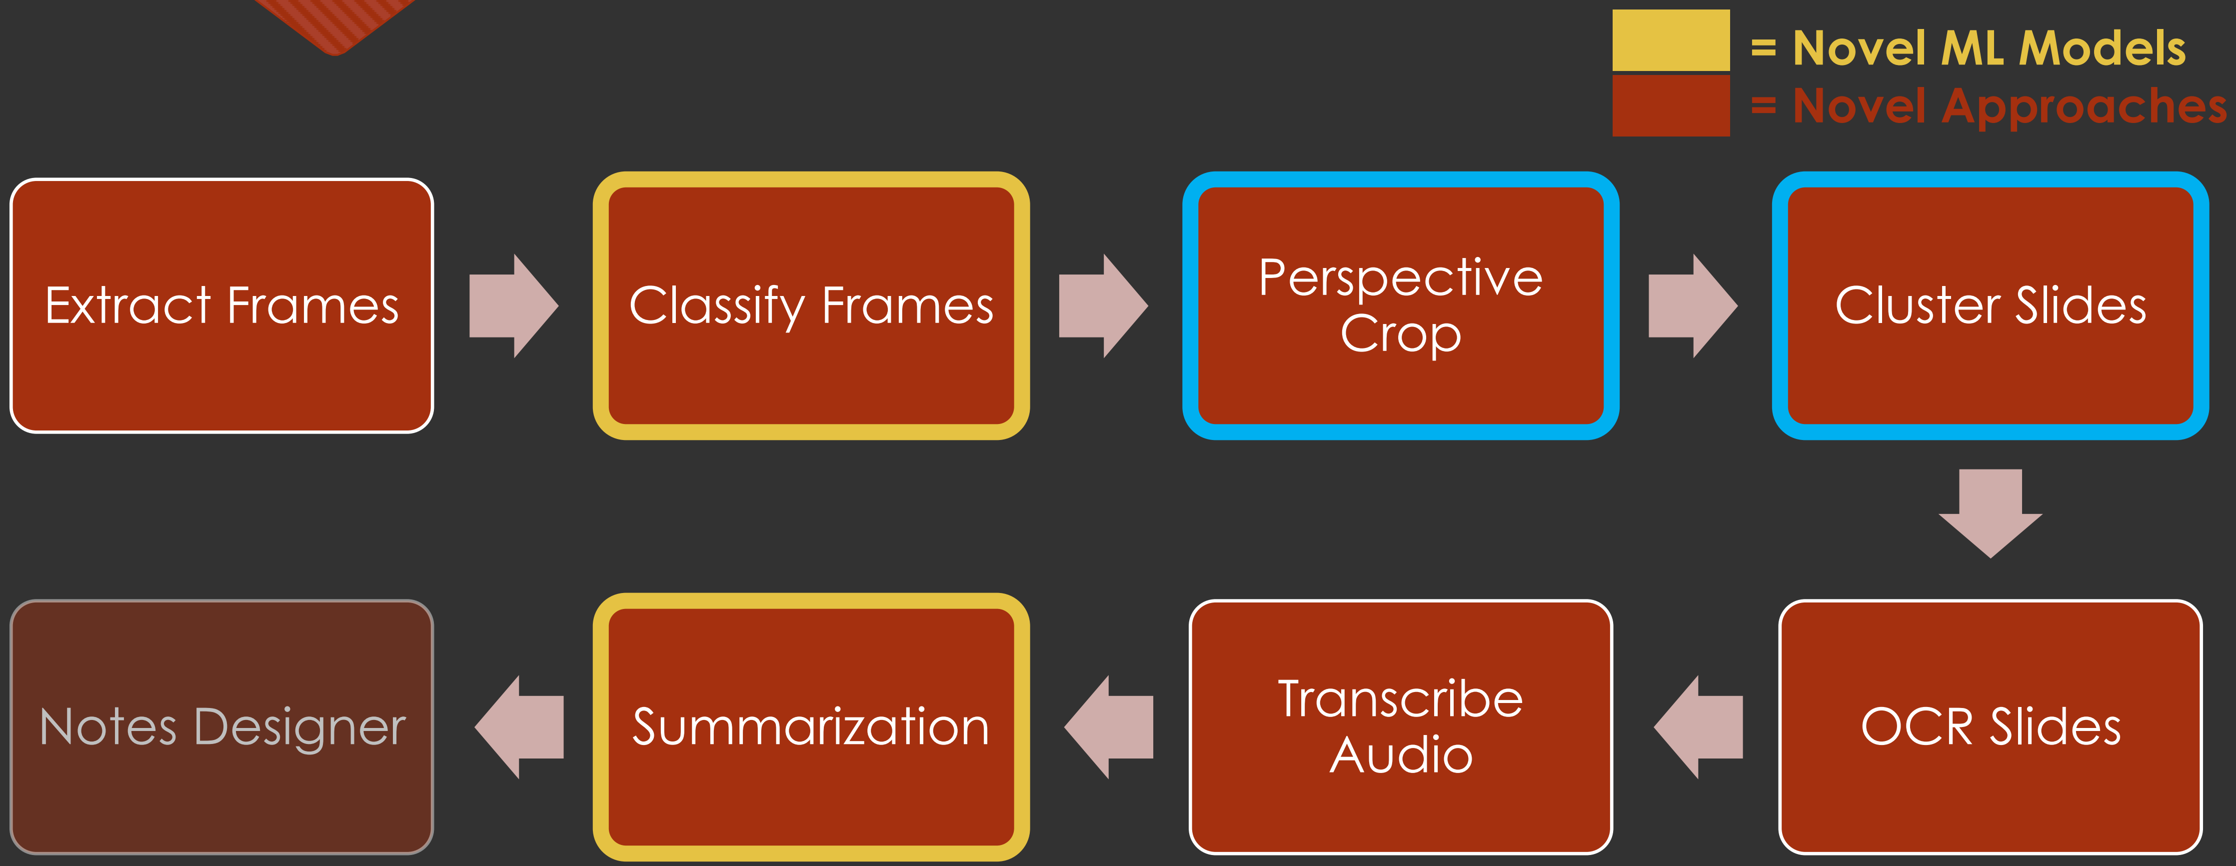

Text Summarization – PreSumm Example

Original Article

visitors to a wildlife park in new zealand got to encounter a pride of lions up-close and personal . filmed at the orana wildlife park -- the country 's only open-range zoo -- the video shows the lions interacting with the visitors who stand inside a metal cage attached to a car . the video , which was captured by ekant veer , 35 , an associate professor at the university of canterbury , also shows the lions scaling the cage and eating meat through its bars . the lions at orana wildlife park approach the metal cage and begin interacting with the people inside standing with its paws against the cage , a lion is introduced as sakura who is around 11-years-old and weighs about 265kg . as the keeper speaks , the lion licks at a piece of meat that is held up against the bars as another lion walks across the roof of the cage . looking down at the people below , the lion wanders around as if deciding who it would like to make its prey before staring down the lens of the camera . one of the lions notices meat and begins sticking out its tongue in the hope of being fed a lion stands next to one of the keepers and its large paw is the same size as the lady 's head the people inside can be seen recording the many lions from their phones , while another -- with paws the same size as the keeper 's head -- holds itself up against the cage and chews on some meat . later in the video people can be seen pointing out the various felines as a keeper moves her hand along the cage , instigating the lion to follow . still frames capture a lion standing up against the side of the cage alongside the keeper -- its power and size is plain to see . a keeper holds a piece of meat up to the bars of the cage and a lion follows her hand in the hope of receiving it a number of lions are fed directly through the metal bars , while others receive meat dropped from the back of the cage the car then begins driving away and the lions can be seen chasing after the people in the hope of receiving more food . a keeper then drops meat from the back of it and the lions begin tailing off one by one with their own little piece of food . the video concludes with one lion picking up a final", "reference": "video shows the lions scaling the cage to look at the people inside <q> lions jump up on the side of the bars and eats meat through them <q> the encounter took place at the orana wildlife park in new zealand

Article Link: http://bit.ly/hhousen-summarization

Summarized Version (Abstractive)

the video was filmed at the orana wildlife park - the country 's only open-range zoo <q> the video shows the lions interacting with the visitors who stand inside a metal cage attached to a car <q> the people inside can be seen recording the many lions from their phones , while another holds itself up against the cage and chews on some meat

Text Summarization

Extractive Summarization

Identifies important sections of the text and generates them verbatim producing a **subset** of the sentences **from** the **original text**

Abstractive Summarization

Reproduces important material in a new way after **interpretation** and **examination** of the **text** using **advanced natural language**

Text Summarization

Approaches to text summarization greatly vary[25]

headlines (from around the world)	outlines (notes for students)	minutes (of a meeting)	previews (of movies)
synopses (soap opera listings)	reviews (of a book, CD, movie, etc.)	digests (TV guide)	biography (resumes, obituaries)
abridgments (Shakespeare for children)	bulletins (weather forecasts/stock market reports)	sound bites (politicians on a current issue)	histories (chronologies of salient events)

Text Summarization

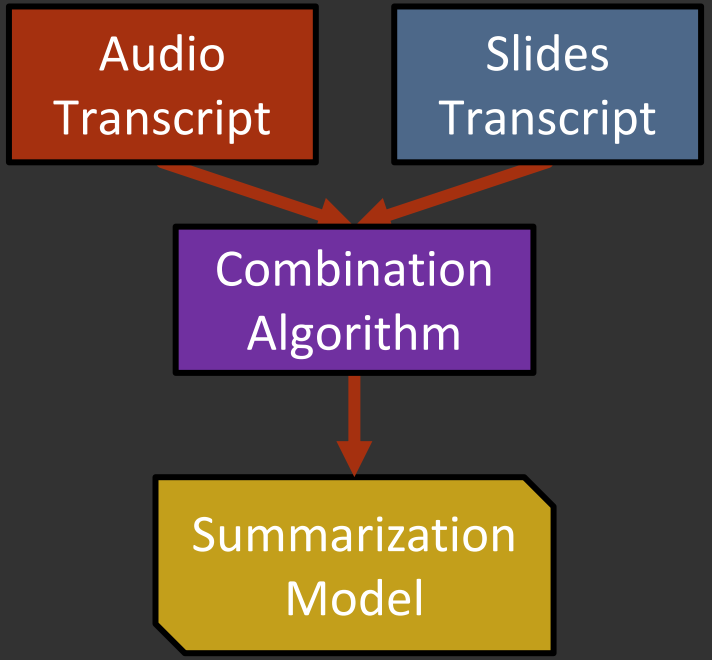

Text Summarization

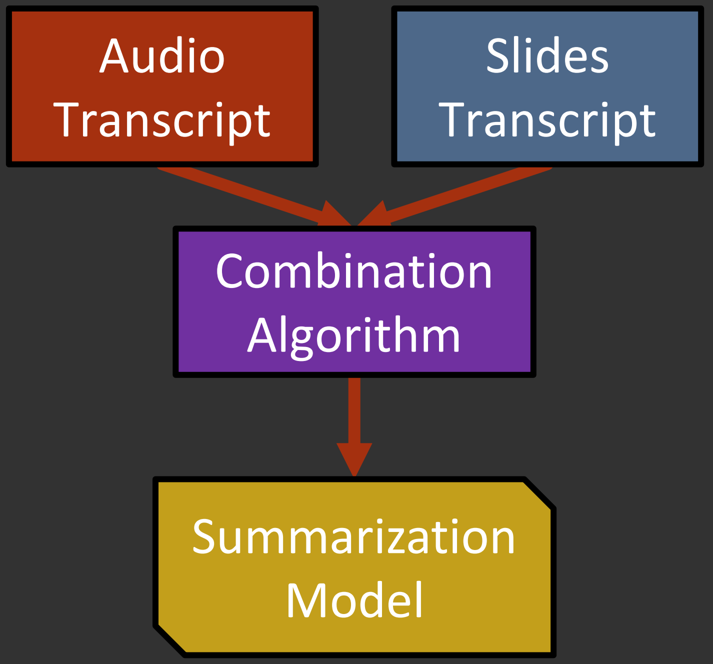

Combination Algorithm

- **only_asr** – only uses the audio transcript (deletes the slide transcript)
- **only_slides** – reverse of only_asr
- **concat** – appends audio transcript to slide transcript
- **full_sents** – audio transcript appended to only the complete sentences of the slide transcript
- **keyword_based** (most advanced) – selects a certain percentage of sentences from the audio transcript based on keywords found in the slides transcript

Text Summarization

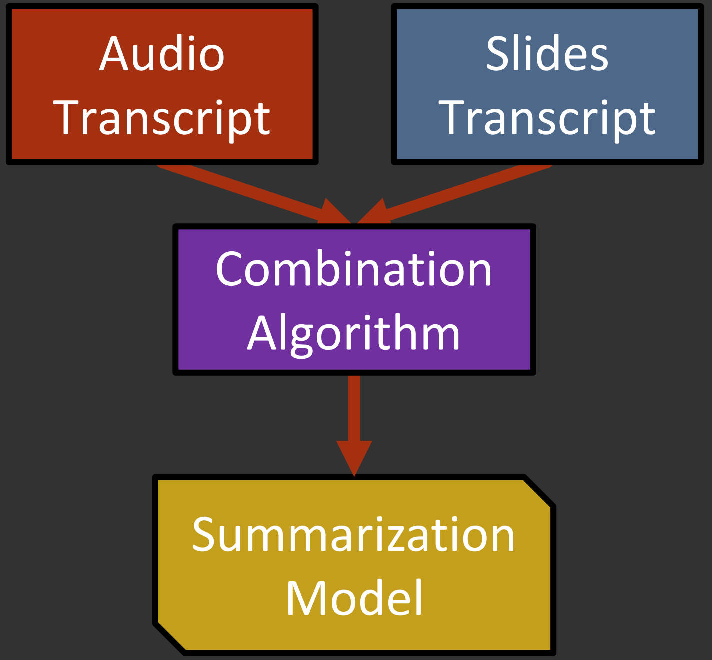

Text Summarization

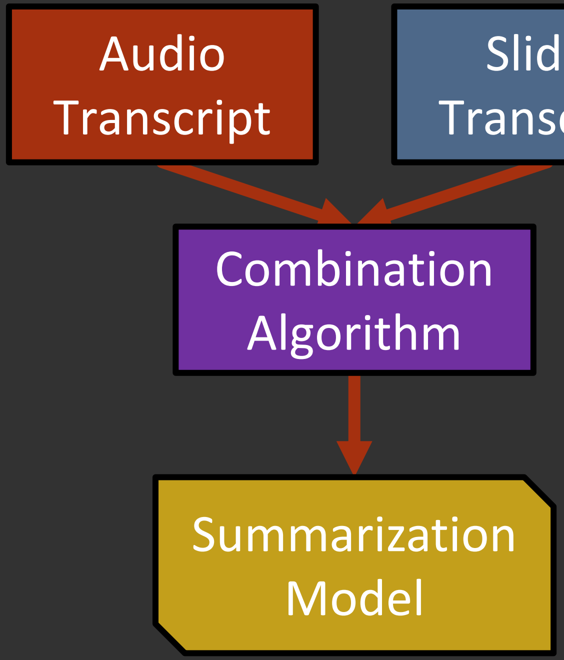

Summarization Model

1. Modifications

- Get only complete sentences

2. Extractive Summarization

- Cluster – groups the lecture transcript into categories by topic and summarizes each topic using "generic"
- Generic (non-neural) – uses algorithms from "sumy" package: lsa, luhn, lex_rank, text_rank, edmundson, random

3. Abstractive Summarization

- PreSumm or BART or TransformerExtSum

Text Summarization – PreSumm

[24]

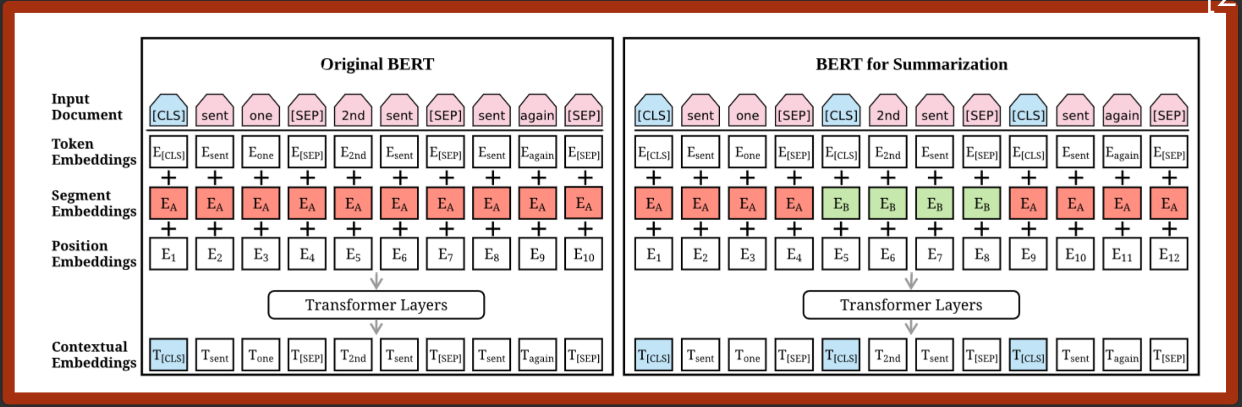

- Paper: "**Text Summarization with Pretrained Encoders**"
- Based on **BERT**
- BERT is **pretrained language model** (understands English language) that can be applied to natural language processing (**NLP**) tasks
- Applied BERT to **text summarization**

Text Summarization – BART

- Paper: “**BART: Denoising Sequence-to-Sequence Pre-training for Natural Language Generation, Translation, and Comprehension**”
- BART is trained by
  1. Corrupting text with an arbitrary noising function
  2. Learning a model to reconstruct the original text
  3. Finetuning to summarize

Text Summarization – Result Comparison

	CNN/DailyMail			XSum		
	R1	R2	RL	R1	R2	RL
Lead-3	40.42	17.62	36.67	16.30	1.60	11.95
PTGEN (See et al., 2017)	36.44	15.66	33.42	29.70	9.21	23.24
PTGEN+COV (See et al., 2017)	39.53	17.28	36.38	28.10	8.02	21.72
UniLM	43.33	20.21	40.51	-	-	-
BERTSUMABS (Liu & Lapata, 2019)	41.72	19.39	38.76	38.76	16.33	31.15
BERTSUMEXTABS (Liu & Lapata, 2019)	42.13	19.60	39.18	38.81	16.50	31.27
BART	**44.16**	**21.28**	**40.90**	**45.14**	**22.27**	**37.25**

BART > PreSumm

Text Summarization – TransformerExtSum

- Being **created** by **HHousen**
- **Improves** the **architecture** of **BertSum** (extractive component of PreSumm)
- Final results nearly completed (waiting for powerful GPU access)

Text Summarization – TransformerExtSum

Name	ROUGE-1	ROUGE-2	ROUGE-L
distilbert-base-uncased **My Improvement**	41.1	18.8	26.5
distilbert-base-uncased **BertSum**	40.1	18.1	26.0
bert-base-uncased **My Improvement**	40.7	18.7	26.6
bert-base-uncased **BertSum**	40.2	18.2	26.1

Pooling mode improvements associated with a **0.617 average** ROUGE F_1 score improvement.

Text Summarization – TransformerExtSum

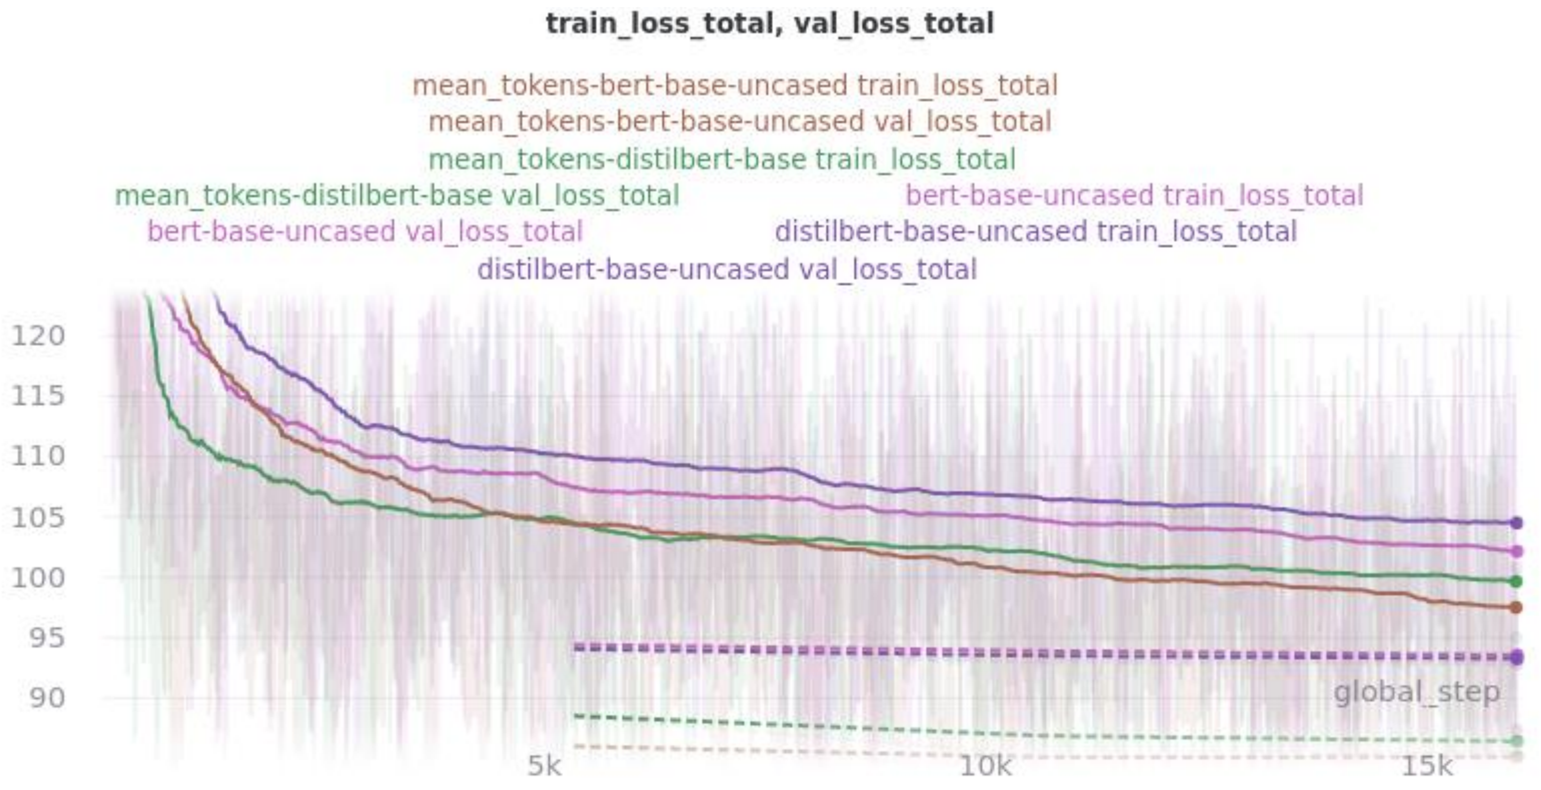

Loss graphs from pooling experiment

3. End-to-End Approach – Diagram

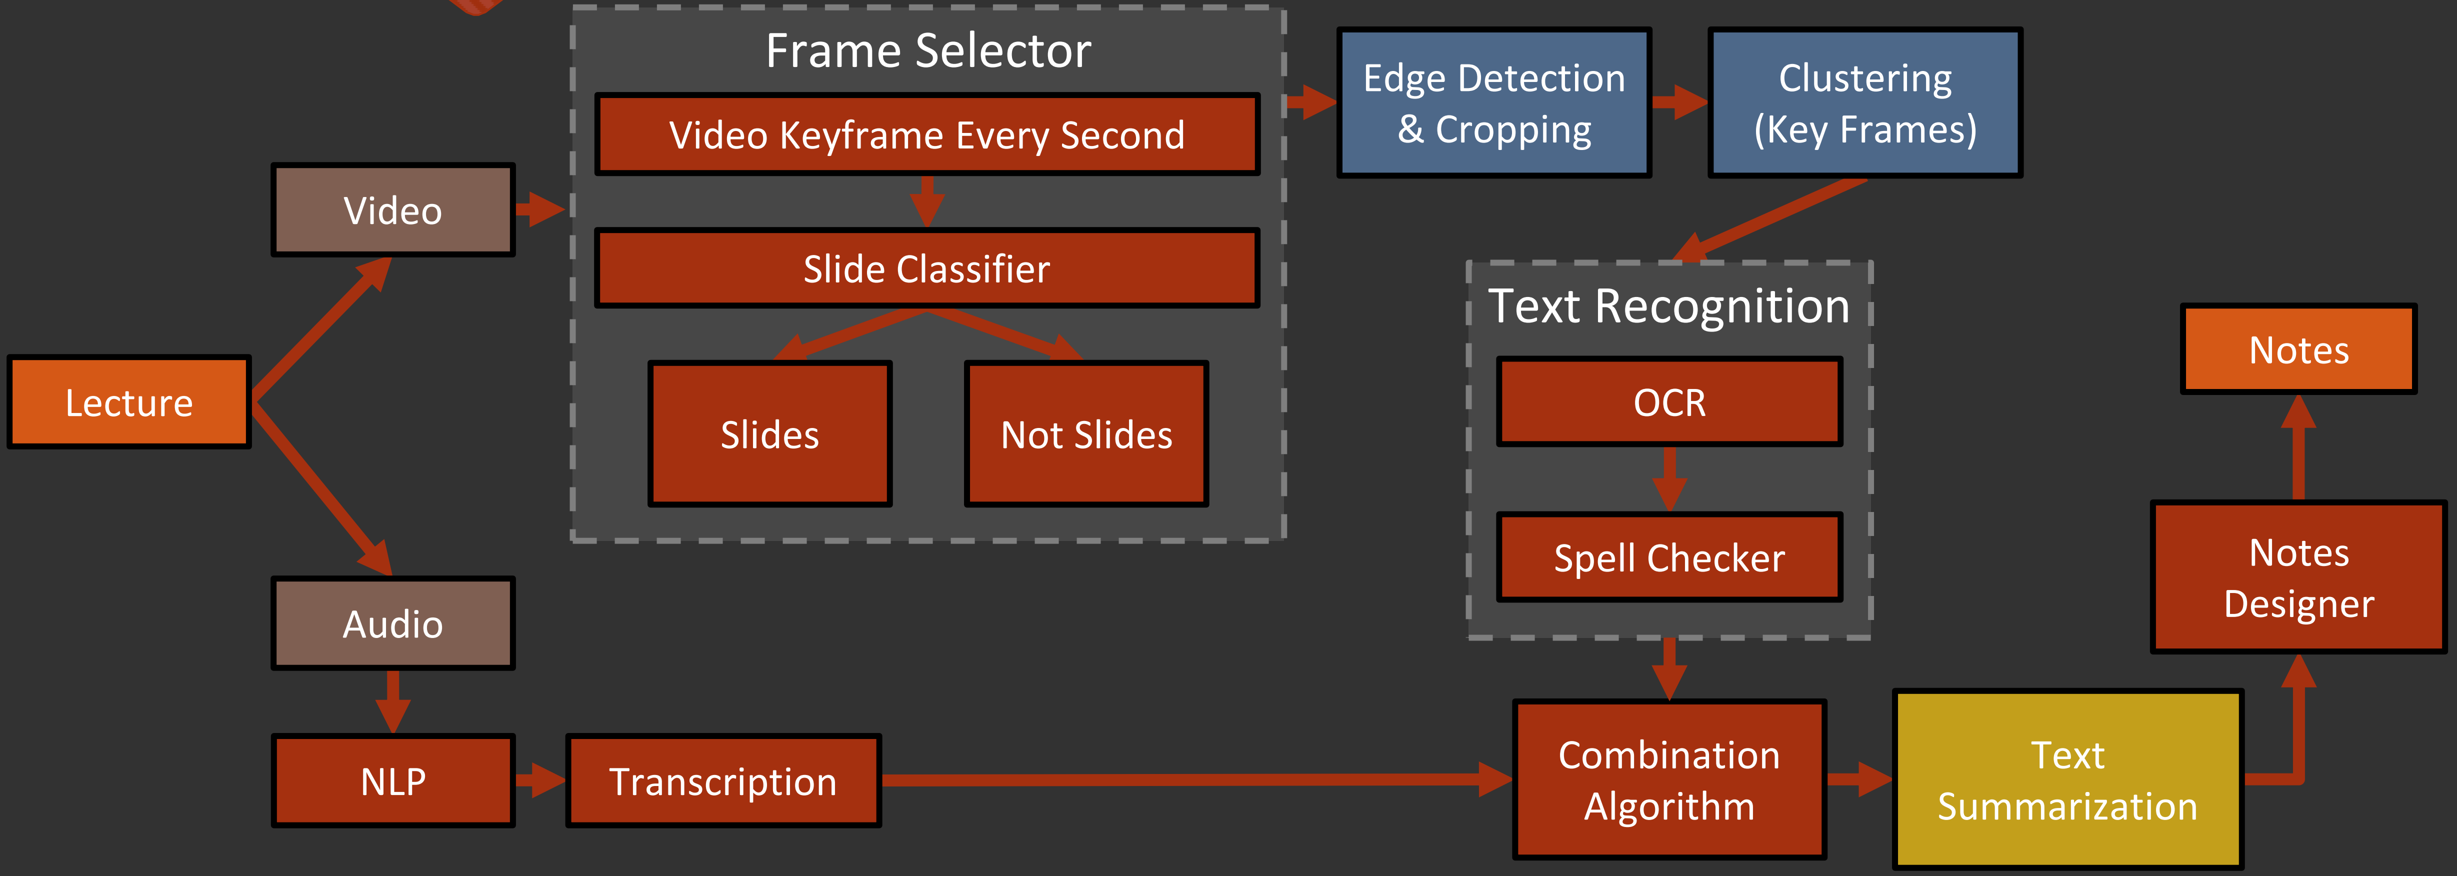

Notes Designer

- **Basic formatting** (bold, italics, etc.) from **OCR**
- **Features specific** to **notes** (headings, etc.) will be generated by the **Notes Designer**
- **Implementation** details **not determined**

4. Website

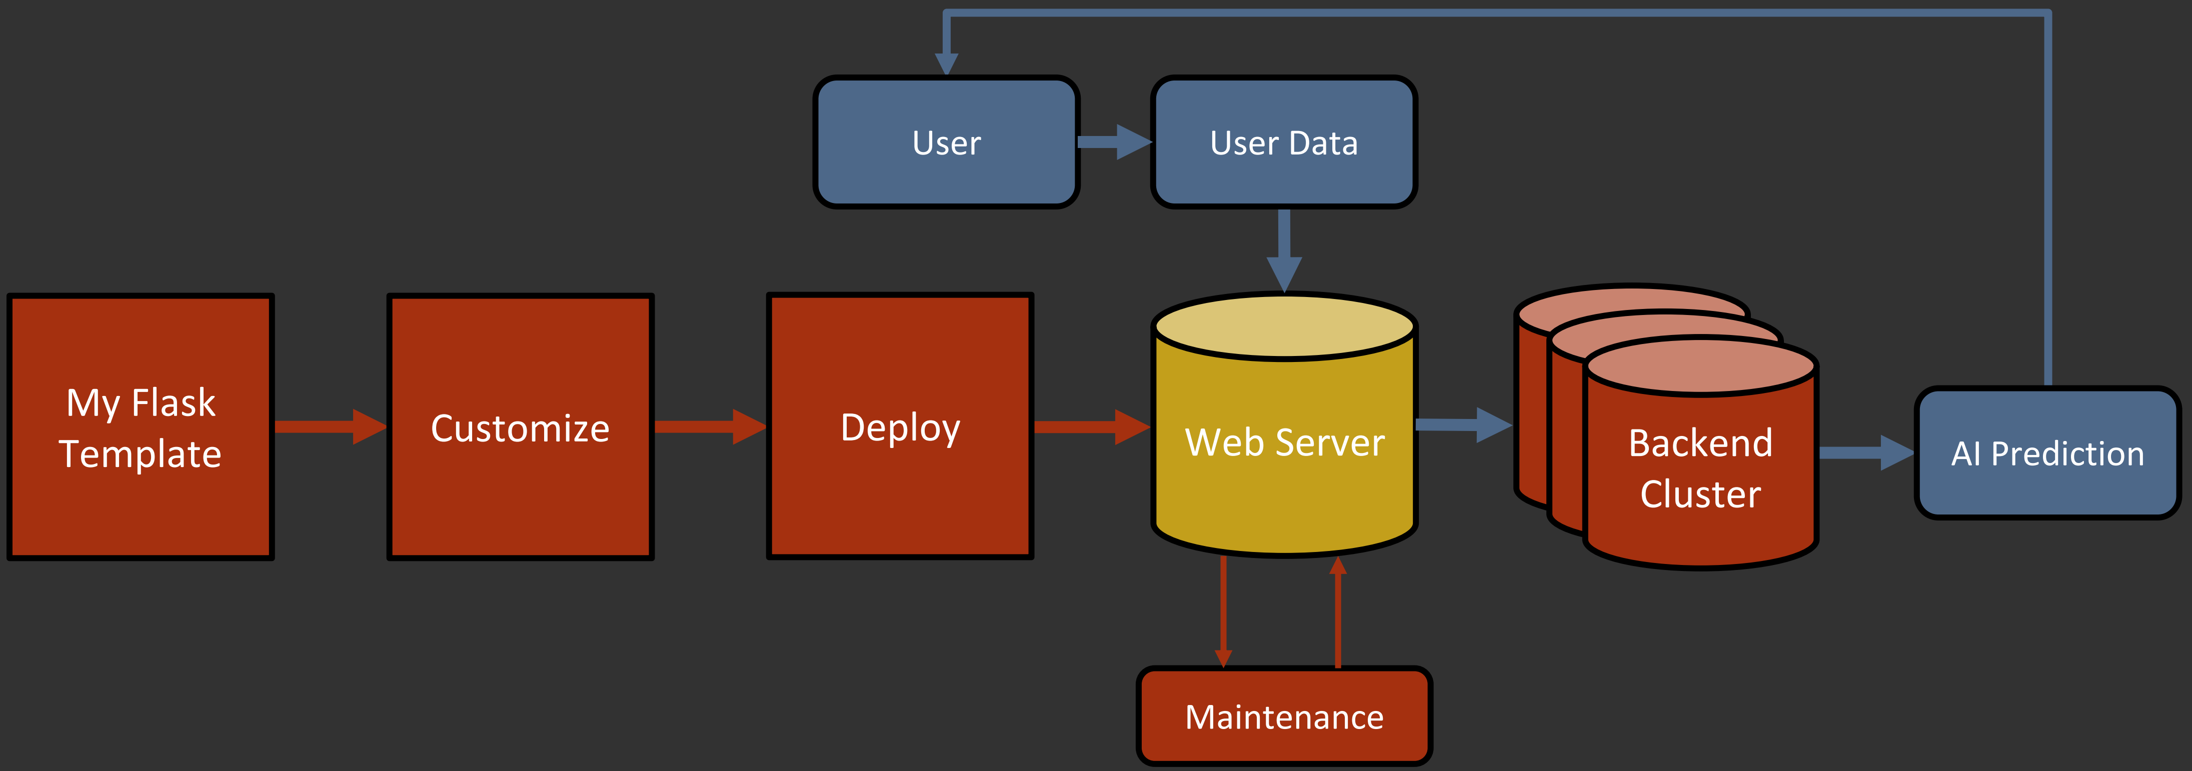
User
User Data
My Flask Template
Customize
Deploy
Web Server
Backend Cluster
AI Prediction
Maintenance

Anticipated Results

https://unsplash.com/photos/OyCl7Y4y0Bk

Education Process

Summaries of lecture slide sets **enhanced likelihood** students would **preview material** & sometimes led to **better quiz scores** [21]

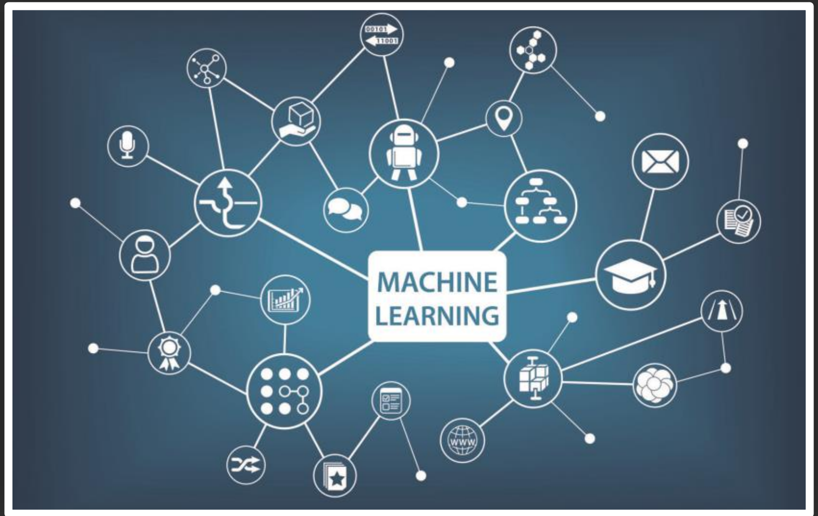

https://www.advectas.com/en/blog/what-is-machine-learning/

Machine Learning

Slide classifier (makes future **slide analysis easier**), **dataset** of lecture videos, text summarization, **OCR**

Future/Next Steps

</>

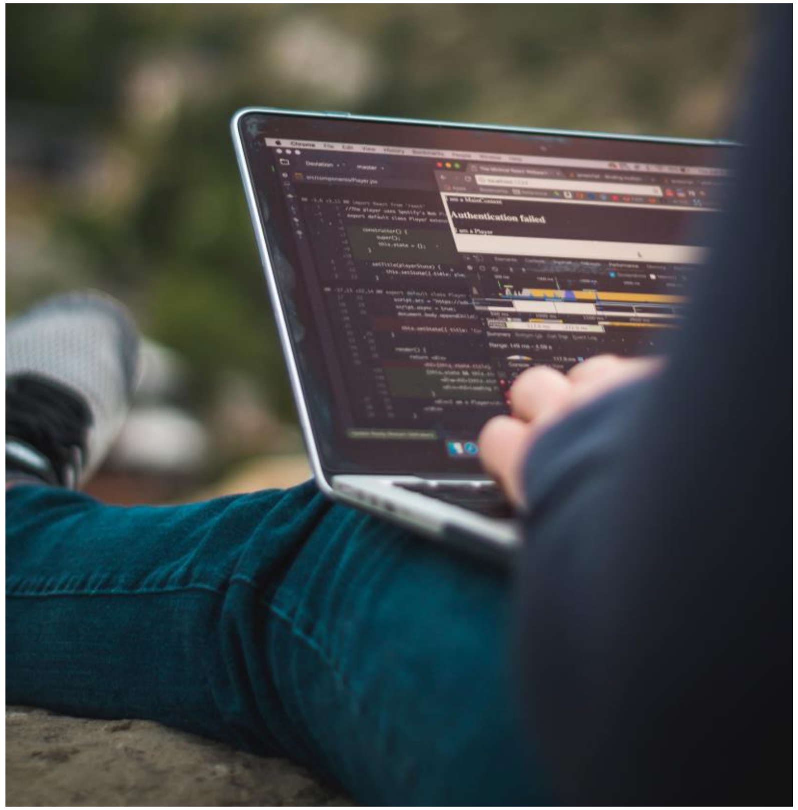

https://unsplash.com/photos/Yhc7YGZlz3g

Continue Writing The Code

https://unsplash.com/photos/zFSo6bnZJTw

Assess Impact on Education

Summary – Five-Fold Contribution

Lecture Video Dataset
- Appropriate categories
- Large variety

Slide Classifier
- Identify important frames in slide presentations

Summarization Models
- Improve state-of-the-art in text summarization
- Novel approaches

End-To-End
- Multimodal approach
- Convert lecture videos to notes

Online Web Service
- Allows people to easily use the created lecture summarizer to create notes based on their own lectures

Acknowledgements

- Dr. Dhiraj Joshi for giving me advice on methodology
- Gillian Rinaldo for guiding me
- Classmates/Peers for supporting me
- Parents for helpful suggestions

References

- [1] Szeliski, R. (2010). *Computer Vision: Algorithms and Applications*. Springer.
- [2] Boden, M. (2008). Mind as machine (p. 781). Oxford: Clarendon Press.
- [3] Dragonette, L. (2018, April 19). Artificial Intelligence - AI. Retrieved December 6, 2018, from https://www.investopedia.com/terms/a/artificial-intelligence-ai.asp
- [4] Machine learning. (2018, December 04). Retrieved from https://en.wikipedia.org/wiki/Machine_learning
- [5] Feature (computer vision). (2019). Retrieved from https://en.wikipedia.org/wiki/Feature_(computer_vision)
- [6] Ripley, B. (2009). Pattern recognition and neural networks (p. 354). Cambridge: Cambridge University Press.
- [7] What is the difference between labeled and unlabeled data?. (2019). Retrieved from https://stackoverflow.com/questions/19170603/what-is-the-difference-between-labeled-and-unlabeled-data/19172720#19172720
- [8] Cortes, C., & Vapnik, V. (1995). Support-vector networks. Machine Learning, 20(3), 273-297. doi: 10.1007/bf00994018
- [9] A Beginner's Guide to Neural Networks and Deep Learning. (2019). Retrieved from https://skymind.ai/wiki/neural-network
- [10] Le, J. (2018). A Tour of The Top 10 Algorithms for Machine Learning Newbies. Retrieved from https://towardsdatascience.com/a-tour-of-the-top-10-algorithms-for-machine-learning-newbies-dde4edffae11
- [11] Brownlee, J. (2017). What is the Difference Between Test and Validation Datasets?. Retrieved from https://machinelearningmastery.com/difference-test-validation-datasets/
- [12] Krizhevsky, A., Sutskever, I., & Hinton, G. E. (2012). ImageNet Classification with Deep Convolutional Neural Networks. *Advances in Neural Information Processing Systems,* 1097-1105. Retrieved February 5, 2019.
- [13] Image Classification. (n.d.). Retrieved February 5, 2019, from https://cs231n.github.io/classification/
- [14] Seif, G. (2018, December 13). How to do everything in Computer Vision. Retrieved February 3, 2019, from https://towardsdatascience.com/how-to-do-everything-in-computer-vision-2b442c469928
- [15] Redmon, J., Divvala, S., Girshick, R., & Farhadi, A. (2016). You Only Look Once: Unified, Real-Time Object Detection. *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. doi:10.1109/cvpr.2016.91
- [16] Chablani, M. (2017, August 21). YOLO - You only look once, real time object detection explained. Retrieved February 5, 2019, from https://towardsdatascience.com/yolo-you-only-look-once-real-time-object-detection-explained-492dc9230006
- [17] Gandhi, R. (2018, July 09). R-CNN, Fast R-CNN, Faster R-CNN, YOLO - Object Detection Algorithms. Retrieved February 5, 2019, from https://towardsdatascience.com/r-cnn-fast-r-cnn-faster-r-cnn-yolo-object-detection-algorithms-36d53571365e

References

- [18] Krizhevsky, A. (n.d.). Google Scholar. Retrieved February 5, 2019, from https://scholar.google.com/citations?user=xegzhJcAAAAJ
- [19] Mahasseni, B., Lam, M., & Todorovic, S. (2017). Unsupervised Video Summarization with Adversarial LSTM Networks. 2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR). doi:10.1109/cvpr.2017.318
- [20] Yang, H., Siebert, M., Luhne, P., Sack, H., & Meinel, C. (2012). Lecture Video Indexing and Analysis Using Video OCR Technology. Journal of Multimedia Processing and Technologies, 2(4), 176-196. doi:10.1109/sitis.2011.20
- [21] Shimada, A., Okubo, F., Yin, C., & Ogata, H. (2017). Automatic Summarization of Lecture Slides for Enhanced Student Preview Technical Report and User Study. IEEE Transactions on Learning Technologies, 11(2), 165-178. doi:10.1109/tlt.2017.2682086
- [22] Adcock, J., Cooper, M., Denoue, L., Pirsiavash, H., & Rowe, L. A. (2010). TalkMiner. MM'10 - Proceedings of the ACM Multimedia 2010 International Conference, 241-250. doi:10.1145/1873951.1873986
- [23] "Cisco Visual Networking Index: Forecast and Trends, 2017–2022 White Paper." VNI Global Fixed and Mobile Internet Traffic Forecasts. Cisco, 27 February 2019, https://www.cisco.com/c/en/us/solutions/collateral/service-provider/visual-networking-index-vni/white-paper-c11-741490.html
- [24] Liu, Y., & Lapata, M. (2019). Text Summarization with Pretrained Encoders. Retrieved from https://arxiv.org/abs/1908.08345
- [25] Sciforce. (2019, January 23). Towards Automatic Text Summarization: Extractive Methods. Retrieved December 27, 2019, from https://medium.com/sciforce/towards-automatic-text-summarization-extractive-methods-e8439cd54715.
- [26] He, K., Zhang, X., Ren, S., & Sun, J. (2016). Deep Residual Learning for Image Recognition. *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. doi: 10.1109/cvpr.2016.90
- [27] Pauline A. Nye, Terence J. Crooks, Melanie Powley, and Gail Tripp. Student note-taking related to university examination performance. Higher Education, 13(1):85–97, February 1984. ISSN 0018-1560, 1573-174X. doi: 10.1007/BF00136532. URL http://link.springer.com/10.1007/BF00136532.
- [28] Kenneth A. Kiewra. Providing the Instructor's Notes: An Effective Addition to Student Notetaking. Educational Psychologist, 20(1):33–39, January 1985. ISSN 0046-1520, 1532-6985. doi: 10.1207/s15326985ep2001_5. URL https://www.tandfonline.com/doi/full/10.1207/s15326985ep2001_5.
- [29] Jennifer L Austin, Matthew D Thibeault, and Jon S Bailey. Effects of Guided Notes on University Students' Responding and Recall of Information. Journal of Behavioral Education, page 12, 2002.
- [30] Anneh Mohammad Gharravi. Impact of instructor-provided notes on the learning and exam performance of medical students in an organ system-based medical curriculum. Advances in Medical Education and Practice, Volume 9:665–672, September 2018. ISSN 1179-7258. doi: 10.2147/AMEP.S172345. URL https://www.dovepress.com/impact-of-instructor-provided-notes-on-the-learning-and-exam-performan-peer-reviewed-article-AMEP.
- [31] Atsushi Shimada, Fumiya Okubo, Chengjiu Yin, and Hiroaki Ogata. Automatic Summarization of Lecture Slides for Enhanced Student PreviewTechnical Report and User Study. IEEE Transactions on Learning Technologies, 11(2):165–178, April 2018. ISSN 1939-1382, 2372-0050. doi: 10.1109/TLT.2017.2682086. URL https://ieeexplore.ieee.org/document/7879355/.

Lecture2notes: Summarizing Lecture Videos by Classifying Slides and Analyzing Text

Hayden Housen
Programming • Computer Vision • Machine Learning

https://haydenhousen.com
hayden@haydenhousen.com
GitHub: @HHousen
Mentor: Dhiraj Joshi @ IMB

All Images Without Citations Are My Own
Background: https://unsplash.com/photos/ieic5Tq8YMk