
EVA Video Database System

Georgia Tech Database Group

Dec 17, 2022

OVERVIEW

1	What is EVA?	3
2	Key Features	5
3	Next Steps	7
4	Illustrative EVA Applications	9
4.1	Traffic Analysis Application using Object Detection Model	9
4.2	MNIST Digit Recognition using Image Classification Model	9
4.3	Movie Analysis Application using Face Detection + Emotion Classification Models	9
5	Community	11
6	Video Database System	13
6.1	Usability and Application Maintainability	13
6.2	GPU Cost and Human Time	13
7	Getting Started	15
7.1	Part 1: Install EVA	15
7.2	Launch EVA server	15
7.3	Part 2: Start a Jupyter Notebook Client	15
7.4	Part 3: Register an user-defined function (UDF)	16
7.5	Part 5: Start a Command Line Client	17
8	List of Notebooks	19
9	Start EVA Server	21
9.1	Launch EVA server	21
10	MNIST TUTORIAL	23
10.1	Start EVA server	23
10.2	Downloading the videos	23
10.3	Upload the video for analysis	24
10.4	Visualize Video	24
10.5	Create an user-defined function (UDF) for analyzing the frames	24
10.6	Run the Image Classification UDF on video	25
10.7	Visualize output of query on the video	25
11	Object Detection Tutorial	27
11.1	Start EVA server	27
11.2	Download the Videos	27

11.3	Load the surveillance videos for analysis	28
11.4	Visualize Video	28
11.5	Register FasterRCNN Object Detection model in an User-Defined Function (UDF) in EVA	28
11.6	Run Object Detector on the video	29
11.7	Visualizing output of the Object Detector on the video	32
11.8	Dropping an User-Defined Function (UDF)	35
12	MOVIE ANALYSIS	37
12.1	Start EVA Server	37
12.2	Video Files	37
12.3	Adding the video file to EVADB for analysis	38
12.4	Visualize Video	38
12.5	Create an user-defined function(UDF) for analyzing the frames	38
12.6	Run the Face Detection UDF on video	39
12.7	Run the Emotion Detection UDF on the outputs of the Face Detection UDF	39
13	Custom Model Tutorial	47
13.1	Start EVA server	47
13.2	Download custom user-defined function (UDF), model, and video	47
13.3	Load video for analysis	48
13.4	Visualize Video	48
13.5	Create GenderCNN and FaceDetector UDFs	49
13.6	Run Face Detector on video	50
13.7	Composing UDFs in a query	50
13.8	Visualize Output	51
14	EVA Query Language Reference	55
14.1	LOAD	55
14.2	SELECT	56
14.3	EXPLAIN	57
14.4	SHOW	57
14.5	CREATE	57
14.6	DROP	58
14.7	INSERT	58
14.8	RENAME	59
15	User-Defined Functions	61
15.1	Part 1: Writing a custom UDF	61
15.2	Part 2: Registering and using the UDF in queries	63
16	Configure GPU	65
17	Contributing	67
17.1	Setting up the Development Environment	67
17.2	Testing	68
17.3	Submitting a Contribution	68
17.4	Code Style	68
17.5	Troubleshooting	69
18	Debugging	71
18.1	Pre-requisites	71
18.2	Set up Debugger for EVA	71
19	Extending EVA	73
19.1	Command Handler	73

19.2	1. Parser	73
19.3	2. Statement To Plan Convertor	74
19.4	3. Plan Generator	76
19.5	4. Plan Executor	77
19.6	Additional Notes	77
20	EVA Release Guide	79
20.1	Part 1: Before You Start	79
20.2	Part 2: Release Steps	79
21	Packaging	83
21.1	Models	83
21.2	Datasets	83

Multimedia Database System – Where SQL meets Deep Learning



WHAT IS EVA?

EVA is a **database system tailored for video analytics** – think PostgreSQL for videos. It supports a SQL-like language for querying videos like:

- examining the “emotion palette” of different actors
- finding gameplays that lead to a touchdown in a football game

EVA comes with a wide range of commonly used computer vision models. It is written in Python, and it is licensed under the Apache license. The source code is available at <https://github.com/georgia-tech-db/eva>.

If you are wondering why you might need a video database system, start with page on [Video Database Systems](#). It describes how EVA lets users easily make use of deep learning models and how they can reduce money spent on inference on large image or video datasets.

The [Getting Started](#) page shows how you can use EVA for different computer vision tasks: image classification, object detection, action recognition, and how you can easily extend EVA to support your custom deep learning model in the form of user-defined functions.

The [User Guides](#) section contains Jupyter Notebooks that demonstrate how to use various features of EVA. Each notebook includes a link to Google Colab, where you can run the code by yourself.

KEY FEATURES

1. With EVA, you can **easily combine SQL and deep learning models to build next-generation database applications**. EVA treats deep learning models as functions similar to traditional SQL functions like SUM().
2. EVA is **extensible by design**. You can write an **user-defined function** (UDF) that wraps arounds your custom deep learning model. In fact, all the built-in models that are included in EVA are written as user-defined functions.
3. EVA comes with a collection of **built-in sampling, caching, and filtering optimizations** inspired by relational database systems. These optimizations help **speed up queries on large datasets and save money spent on model inference**.

NEXT STEPS

Getting Started A step-by-step guide to installing EVA and running queries

Query Language List of all the query commands supported by EVA

User Defined Functions A step-by-step tour of registering a user defined function that wraps around a custom deep learning model

ILLUSTRATIVE EVA APPLICATIONS

4.1 Traffic Analysis Application using Object Detection Model

4.2 MNIST Digit Recognition using Image Classification Model

4.3 Movie Analysis Application using Face Detection + Emotion Classification Models

COMMUNITY

Join the EVA community on [Slack](#) to ask questions and to share your ideas for improving EVA.



JOIN THE SLACK CHANNEL

Talk to our developers and meet
other members of our community.



VIDEO DATABASE SYSTEM

Over the last decade, deep learning models have radically changed the world of computer vision. They are accurate on a variety of tasks ranging from image classification to emotion detection. However, there are two challenges that prevent a lot of users from benefiting from these models.

6.1 Usability and Application Maintainability

To use a vision model, the user must do a lot of imperative programming across low-level libraries, like OpenCV and PyTorch. This is a tedious process that often leads to a complex program or Jupyter Notebook that glues together these libraries to accomplish the given task. This programming complexity **prevents a lot of people who are experts in other domains from benefiting from these models**.

Historically, database systems have been successful because the **query language is simple enough** in its basic structure that users without prior experience are able to learn a usable subset of the language on their first sitting. EVA supports a declarative SQL-like query language, called EVAQL, that is designed to make it easier for users to leverage these models. With this query language, the user may **compose multiple models in a single query** to accomplish complicated tasks with **minimal programming**.

Here is an illustrative query that examines the emotions of actors in a movie by leveraging multiple deep learning models that take care of detecting faces and analyzing the emotions of the detected bounding boxes:

```
SELECT id, bbox, EmotionDetector(Crop(data, bbox))
FROM Interstellar
      JOIN LATERAL UNNEST(FaceDetector(data)) AS Face(bbox, conf)
WHERE id < 15;
```

By using a declarative language, the complexity of the program or Jupyter Notebook is significantly reduced. This in turn leads to more **maintainable code** that allows users to build on top of each other's queries.

6.2 GPU Cost and Human Time

From a cost standpoint, it is very expensive to run these deep learning models on large image or video datasets. For example, the state-of-the-art object detection model takes multiple GPU-decades to process just a year's worth of videos from a single traffic monitoring camera. Besides the money spent on hardware, this also increases the time that the user spends waiting for the model inference process to finish.

EVA **automatically** optimizes the queries to **reduce inference cost and query execution time** using its Cascades-style query optimizer. EVA's optimizer is tailored for video analytics. The **Cascades-style extensible query optimization framework** has worked very well for several decades in SQL database systems. Query optimization is one of the signature components of database systems — **the bridge that connects the declarative query language to efficient execution**.

GETTING STARTED

7.1 Part 1: Install EVA

EVA supports Python (versions ≥ 3.7). To install EVA, we recommend using the pip package manager:

```
pip install evadb
```

7.2 Launch EVA server

EVA is based on a [client-server architecture](#). To launch the EVA server, run the following command on the terminal:

```
eva_server &
```

7.3 Part 2: Start a Jupyter Notebook Client

Here is an [illustrative Jupyter notebook](#) focusing on MNIST image classification using EVA. The notebook works on [Google Colab](#).

7.3.1 Connect to the EVA server

To connect to the EVA server in the notebook, use the following Python code:

```
# allow nested asyncio calls for client to connect with server
import nest_asyncio
nest_asyncio.apply()

# hostname and port of the server where EVA is running
connection = connect(host = '0.0.0.0', port = 5432)

# cursor allows the notebook client to send queries to the server
cursor = connection.cursor()
```

7.3.2 Load video for analysis

Download the MNIST video.

```
!wget -nc https://www.dropbox.com/s/yxljxz6zxoqu54v/mnist.mp4
```

Use the LOAD statement is used to load a video onto a table in EVA server.

```
cursor.execute('LOAD FILE "mnist.mp4" INTO MNISTVideoTable;')
response = cursor.fetch_all()
print(response)
```

7.3.3 Run a query

Run a query over the video to retrieve the output of the MNIST CNN function that is included in EVA as a built-in user-defined function (UDF).

```
cursor.execute("""SELECT id, MnistCNN(data).label
                  FROM MNISTVideoTable
                  WHERE id < 5;""")
response = cursor.fetch_all()
print(response)
```

That's it! You can now run more complex queries.

7.4 Part 3: Register an user-defined function (UDF)

User-defined functions allow us to combine SQL with deep learning models. These functions can wrap around deep learning models.

Download an user-defined function for classifying MNIST images.

```
!wget -nc https://raw.githubusercontent.com/georgia-tech-db/eva/master/tutorials/apps/
↪mnist/eva_mnist_udf.py
```

```
cursor.execute("""CREATE UDF IF NOT EXISTS MnistCNN
                  INPUT  (data NDARRAY (3, 28, 28))
                  OUTPUT (label TEXT(2))
                  TYPE    Classification
                  IMPL    'eva_mnist_udf.py';
                  """)
response = cursor.fetch_all()
print(response)
```

7.4.1 Run a more interesting query using the newly registered UDF

```
cursor.execute("""SELECT data, MnistCNN(data).label
                  FROM MNISTVideoTable
                  WHERE id = 30;""")
response = cursor.fetch_all()
```

7.4.2 Visualize the Output

The output of the query is [visualized in the notebook](#).

7.5 Part 5: Start a Command Line Client

Besides the notebook interface, EVA also exports a command line interface for querying the server. This interface allows for quick querying from the terminal:

```
>>> eva_client
eva=# LOAD FILE "mnist.mp4" INTO MNISTVid;
@status: ResponseStatus.SUCCESS
@batch:

0 Video successfully added at location: mnist.p4
@query_time: 0.045

eva=# SELECT id, data FROM MNISTVid WHERE id < 1000;
@status: ResponseStatus.SUCCESS
@batch:
```

	mnistvid.id	mnistvid.data
0	0	[[[0 2 0]\n [0 0 0]\n...
1	1	[[[2 2 0]\n [1 1 0]\n...
2	2	[[[2 2 0]\n [1 2 2]\n...
..	...	
997	997	[[[0 2 0]\n [0 0 0]\n...
998	998	[[[0 2 0]\n [0 0 0]\n...
999	999	[[[2 2 0]\n [1 1 0]\n...

```
[1000 rows x 2 columns]
@query_time: 0.216

eva=# exit
```


LIST OF NOTEBOOKS

This section contains Jupyter Notebooks that demonstrate how to use various features of EVA. Each notebook includes a link to Google Colab, where you can run the code by yourself.

- [Starting EVA Server](#)
- [MNIST Image Classification](#)
- [Object Detection](#)
- [Emotion Analysis](#)
- [Using a Custom Model](#)

START EVA SERVER

9.1 Launch EVA server

We use this notebook for launching the EVA server.

```
## Install EVA package if needed
##pip install "evadb[udf]" --quiet

import os
import time

def shell(command):
    print(command)
    os.system(command)

def stop_eva_server():
    # Kill any process listening on EVA's standard port
    shell('[ -z "$(lsof -ti:5432)" ] || kill -9 "$(lsof -ti:5432)"')

def launch_eva_server():
    # Stop EVA server if it is running
    stop_eva_server()

    os.environ['GPU_DEVICES'] = '0'

    # Start EVA server
    shell("nohup eva_server > eva.log 2>&1 &")

    try:
        with open('eva.log', 'r') as f:
            print(f.read())
    except FileNotFoundError:
        pass

    # Wait for server to start
    time.sleep(10)

def connect_to_server():
    from eva.server.db_api import connect
    %pip install nest_asyncio --quiet
```

(continues on next page)

(continued from previous page)

```
import nest_asyncio
nest_asyncio.apply()

# Connect client with server
connection = connect(host = '127.0.0.1', port = 5432)
cursor = connection.cursor()

return cursor

# Launch server
launch_eva_server()
```

```
[ -z "$(lsof -ti:5432)" ] || kill -9 "$(lsof -ti:5432)"
```

```
nohup eva_server > eva.log 2>&1 &
```

MNIST TUTORIAL

10.1 Start EVA server

We are reusing the start server notebook for launching the EVA server.

```
!wget -nc "https://raw.githubusercontent.com/georgia-tech-db/eva/master/tutorials/00-  
↪start-eva-server.ipynb"  
%run 00-start-eva-server.ipynb  
cursor = connect_to_server()
```

```
File '00-start-eva-server.ipynb' already there; not retrieving.
```

```
[ -z "$(lsof -ti:5432)" ] || kill -9 "$(lsof -ti:5432)"
```

```
nohup eva_server > eva.log 2>&1 &
```

```
Note: you may need to restart the kernel to use updated packages.
```

10.2 Downloading the videos

```
# Getting MNIST as a video  
!wget -nc https://www.dropbox.com/s/yxljxz6zxoqu54v/mnist.mp4  
# Getting a udf  
!wget -nc https://raw.githubusercontent.com/georgia-tech-db/eva/master/tutorials/apps/  
↪mnist/eva_mnist_udf.py
```

```
File 'mnist.mp4' already there; not retrieving.
```

```
File 'eva_mnist_udf.py' already there; not retrieving.
```

10.3 Upload the video for analysis

```
cursor.execute('DROP TABLE MNISTVid')
response = cursor.fetch_all()
print(response)
cursor.execute('LOAD FILE "mnist.mp4" INTO MNISTVid')
response = cursor.fetch_all()
print(response)
```

```
@status: ResponseStatus.SUCCESS
@batch:
0
0 Table Successfully dropped: MNISTVid
@query_time: 0.03111536242067814
@status: ResponseStatus.SUCCESS
@batch:
0
0 Video successfully added at location: mnist.mp4
@query_time: 0.043151989579200745
```

10.4 Visualize Video

```
from IPython.display import Video
Video("mnist.mp4", embed=True)
```

```
<IPython.core.display.Video object>
```

10.5 Create an user-defined function (UDF) for analyzing the frames

```
cursor.execute("""CREATE UDF IF NOT EXISTS MnistCNN
                INPUT  (data NDARRAY (3, 28, 28))
                OUTPUT (label TEXT(2))
                TYPE   Classification
                IMPL    'eva_mnist_udf.py';
            """)
response = cursor.fetch_all()
print(response)
```

```
@status: ResponseStatus.SUCCESS
@batch:
0
0 UDF MnistCNN already exists, nothing added.
@query_time: 0.02193857543170452
```




OBJECT DETECTION TUTORIAL

11.1 Start EVA server

We are reusing the start server notebook for launching the EVA server.

```
!wget -nc "https://raw.githubusercontent.com/georgia-tech-db/eva/master/tutorials/00-  
start-eva-server.ipynb"  
%run 00-start-eva-server.ipynb  
cursor = connect_to_server()
```

```
File '00-start-eva-server.ipynb' already there; not retrieving.
```

```
[ -z "$(lsof -ti:5432)" ] || kill -9 "$(lsof -ti:5432)"
```

```
nohup eva_server > eva.log 2>&1 &
```

```
Note: you may need to restart the kernel to use updated packages.
```

11.2 Download the Videos

```
# Getting the video files  
!wget -nc https://www.dropbox.com/s/k00wge9exwkfxz6/ua_detrac.mp4  
# Getting the FastRCNN object detector  
!wget -nc https://raw.githubusercontent.com/georgia-tech-db/eva/master/eva/udfs/fastrcnn_  
object_detector.py
```

```
File 'ua_detrac.mp4' already there; not retrieving.
```

```
File 'fastrcnn_object_detector.py' already there; not retrieving.
```

11.3 Load the surveillance videos for analysis

```
cursor.execute('DROP TABLE ObjectDetectionVideos')
response = cursor.fetch_all()
print(response)
cursor.execute('LOAD FILE "ua_detrac.mp4" INTO ObjectDetectionVideos;')
response = cursor.fetch_all()
print(response)
```

```
@status: ResponseStatus.SUCCESS
@batch:
0
0 Table Successfully dropped: ObjectDetectionVideos
@query_time: 0.03250785917043686
@status: ResponseStatus.SUCCESS
@batch:
0
0 Video successfully added at location: ua_detrac.mp4
@query_time: 0.04492253065109253
```

11.4 Visualize Video

```
from IPython.display import Video
Video("ua_detrac.mp4", embed=True)
```

```
<IPython.core.display.Video object>
```

11.5 Register FasterRCNN Object Detection model in an User-Defined Function (UDF) in EVA

```
cursor.execute("""CREATE UDF IF NOT EXISTS FastRCNNObjectDetector
    INPUT (frame NDARRAY UINT8(3, ANYDIM, ANYDIM))
    OUTPUT (labels NDARRAY STR(ANYDIM), bboxes NDARRAY FLOAT32(ANYDIM, 4),
            scores NDARRAY FLOAT32(ANYDIM))
    TYPE Classification
    IMPL 'fastrcnn_object_detector.py';
""")
response = cursor.fetch_all()
print(response)
```

```
@status: ResponseStatus.SUCCESS
@batch:
0
0 UDF FastRCNNObjectDetector already exists, nothing added.
@query_time: 0.02033083327114582
```

11.6 Run Object Detector on the video

```
cursor.execute("""SELECT id, FastRCNNObjectDetector(data)
                  FROM ObjectDetectionVideos
                  WHERE id < 20""")
response = cursor.fetch_all()
print(response)
```

```
@status: ResponseStatus.SUCCESS
```

```
@batch:
```

```
    objectdetectionvideos.id \
0          0
1          1
2          2
3          3
4          4
5          5
6          6
7          7
8          8
9          9
10         10
11         11
12         12
13         13
14         14
15         15
16         16
17         17
18         18
19         19
```

```
    ↳ fastrcnnobjectdetector.labels \
```

```
0  [person, car, car, car, car, car, car, car, car, car, car, car, car, person, car, car, car, ↳
↳ car, car, car, ...
1  [person, car, car, car, car, car, car, car, car, car, car, car, person, car, car, car, car, ↳
↳ car, car, car, ...
2  [person, car, car, car, car, car, car, car, car, car, car, car, car, person, car, car, car, ↳
↳ car, car, car, ...
3  [person, car, car, car, car, car, car, car, car, car, car, car, car, car, car, person, car, ↳
↳ car, car, car, ...
4  [person, car, car, car, car, car, car, car, car, car, car, car, car, car, car, car, person, ↳
↳ motorcycle, ca...
5  [person, car, car, car, car, car, car, car, car, car, car, car, motorcycle, car, car, person, ↳
↳ car, car, ca...
6  [person, car, car, car, car, car, car, car, car, car, car, car, car, car, motorcycle, person, ↳
↳ car, car, ca...
7  [person, car, car, car, car, car, car, car, car, car, car, car, car, car, motorcycle, person, ↳
↳ car, car, ca...
8  [person, car, car, car, car, car, car, car, car, car, car, car, car, car, motorcycle, ↳
↳ person, car, ca...
```

(continues on next page)

(continued from previous page)

```

9  [person, car, car, car, car, car, car, car, car, car, car, motorcycle, car, car, car, car,
  ↪ car, car, ...
10 [person, car, car, car, car, car, car, car, car, car, car, car, motorcycle, car, car, car, car,
  ↪ person, ca...
11 [person, car, car, car, car, car, car, car, car, car, car, car, car, car, motorcycle, ↪
  ↪ motorcycle, car, car...
12 [person, car, car, car, car, car, car, car, car, car, car, motorcycle, car, car, car, car, car,
  ↪ car, car, ...
13 [person, car, car, car, car, car, car, car, car, motorcycle, car, car, car, car, car, car, car,
  ↪ car, car, ...
14 [person, car, car, car, car, car, car, car, car, car, car, motorcycle, car, car, car, car, car,
  ↪ car, car, ...
15 [person, car, car, car, car, car, car, car, car, car, car, car, motorcycle, car, car, car, car,
  ↪ car, car, ...
16 [person, car, car, car, car, car, car, car, car, car, car, car, car, car, car, car, car, car, ↪
  ↪ motorcycle, ...
17 [person, car, car, car, car, car, car, car, car, car, car, car, car, car, car, car, car, car, ↪
  ↪ motorcycle, ...
18 [person, car, car, car, car, car, car, car, car, car, car, car, car, car, car, car, car, bicycle, ↪
  ↪ car, car, car...
19 [person, car, car, car, car, car, car, car, car, car, car, car, car, car, car, car, car, car, ↪
  ↪ car, car, ca...

↪ fastrcnnobjectdetector.bboxes \
0  [[636.2461 363.62677 670.9286 421.68674], [795.974 198.13597 905.6497 239.
  ↪ 72404], [926.5628...
1  [[636.3488 364.41104 671.457 425.17358], [797.5898 198.4674 909.43506 239.6631 ↪
  ↪ ], [926.5780...
2  [[636.69507 367.78067 671.4994 425.38022], [609.8011 215.57028 726.95276 277.
  ↪ 29547], [799.7043...
3  [[637.3386 370.06775 671.1379 430.22446], [926.5658 108.94128 960. 130.
  ↪ 84044], [841.2449...
4  [[637.39386 371.9589 671.4067 429.7561 ], [844.38947 235.53317 960. 292.
  ↪ 27747], [926.5494...
5  [[637.3752 374.44894 671.94543 431.65906], [846.1043 235.29305 959.76 292.
  ↪ 48502], [926.5932...
6  [[637.50885 376.1439 672.65 435.27576], [851.6319 283.0303 960. 353.
  ↪ 44882], [848.5116...
7  [[636.91095 379.39893 672.60986 441.0763 ], [855.6536 284.40753 960. 353.
  ↪ 15482], [926.5692...
8  [[637.67236 381.3916 673.5471 443.23373], [622.0833 218.62851 740.9644 284.6785 ↪
  ↪ ], [859.5268...
9  [[637.6342 384.85025 673.79156 443.4543 ], [851.62354 235.61482 960. 289.6144 ↪
  ↪ ], [926.5669...
10 [[638.32996 386.24515 675.5105 449.4082 ], [853.97705 235.86105 960. 290.
  ↪ 53348], [926.6726...
11 [[638.0419 387.4773 675.44586 452.0051 ], [926.71765 109.181984 960. 130.
  ↪ 68192 ], [856....
12 [[638.6666 389.93994 675.5014 455.6136 ], [858.50885 238.62434 960. 290.6553 ↪
  ↪ ], [926.8178...
13 [[638.99316 393.1623 675.5218 459.98587], [862.4889 237.9856 959.6673 289.5376], ↪
  ↪ [926.9093 10...

```

(continues on next page)

(continued from previous page)

```

14 [[639.02594 395.4505 675.83563 458.40024], [864.2598 237.59782 959.9527 290.
→08386], [926.9506...
15 [[639.6993 397.50488 676.6463 458.95874], [866.5562 239.41269 959.6509 290.
→25293], [890.41 ...
16 [[639.6608 398.67078 677.3171 461.35507], [867.6135 239.08667 959.7032 292.
→32498], [643.4708...
17 [[639.988 401.79584 676.39813 464.388 ], [927.07587 109.3211 959.9775 130.
→52295], [642.4912...
18 [[639.6279 404.75647 677.6399 468.1819 ], [171.9567 113.68296 223.0494 137.
→04474], [872.4612...
19 [[639.65326 407.688 677.55304 471.05777], [170.65564 114.0763 221.7475 137.
→03217], [646.9618...

→fastrcnnobjectdetector.scores
0 [0.9973132, 0.9954666, 0.99453676, 0.9928416, 0.9928219, 0.99168164, 0.9894707, 0.
→98463666, 0.97...
1 [0.9986808, 0.99661547, 0.9946049, 0.9925746, 0.9925718, 0.9904755, 0.9891768, 0.
→9887982, 0.9882...
2 [0.99889356, 0.9963032, 0.99464333, 0.9944518, 0.99270344, 0.9906961, 0.99067444, 0.
→98912454, 0....
3 [0.99847347, 0.99472106, 0.9937597, 0.9915018, 0.99106944, 0.9899698, 0.98859274, 0.
→98738176, 0....
4 [0.99799746, 0.99536246, 0.9944917, 0.9930253, 0.9925647, 0.9924131, 0.9900523, 0.
→98391956, 0.98...
5 [0.99744403, 0.9945592, 0.9944694, 0.9940614, 0.9938199, 0.99164, 0.9875629, 0.
→9839294, 0.982379...
6 [0.99849856, 0.99723524, 0.9956939, 0.99471045, 0.99313545, 0.9920499, 0.99203074, 0.
→988704, 0.9...
7 [0.9991289, 0.99689096, 0.9950283, 0.99449337, 0.99406683, 0.99135935, 0.98904794, 0.
→98687, 0.98...
8 [0.99855894, 0.99789375, 0.99742484, 0.9951934, 0.99357045, 0.99098235, 0.98661834,
→0.9811693, 0...
9 [0.99793184, 0.9957092, 0.99510276, 0.99408025, 0.9926242, 0.992222, 0.98884636, 0.
→98143584, 0.9...
10 [0.9983885, 0.9960192, 0.99516207, 0.9947366, 0.99309456, 0.98936135, 0.9875754, 0.
→9850686, 0.98...
11 [0.99706155, 0.9951617, 0.99343115, 0.9920042, 0.9912471, 0.99058783, 0.98849005, 0.
→9863391, 0.9...
12 [0.99758244, 0.9971852, 0.99477625, 0.99245197, 0.99241346, 0.99238765, 0.9909242, 0.
→9853243, 0....
13 [0.998033, 0.9956489, 0.99477553, 0.9936738, 0.9918057, 0.99007344, 0.9884399, 0.
→9858212, 0.9848...
14 [0.99647397, 0.99619615, 0.9947871, 0.9924697, 0.9908516, 0.98961943, 0.9889376, 0.
→986597, 0.984...
15 [0.99818736, 0.99694735, 0.99624455, 0.99479103, 0.99371606, 0.99213445, 0.9918663,
→0.9894098, 0...
16 [0.99894565, 0.99597675, 0.9959545, 0.99492985, 0.9949143, 0.9911158, 0.9899811, 0.
→98385, 0.9678...
17 [0.998359, 0.9952448, 0.9951807, 0.994565, 0.99443567, 0.99266374, 0.9913652, 0.
→9903177, 0.98383...
18 [0.99799997, 0.9964407, 0.99579465, 0.9955603, 0.99490297, 0.99179196, 0.9901037, 0.
→9847661, 0.9...

```

(continues on next page)

(continued from previous page)

```

19 [0.9981029, 0.9973521, 0.99511445, 0.9949681, 0.9949438, 0.99476826, 0.99410397, 0.
↪9827383, 0.97...
@query_time: 6.682086415588856

```

11.7 Visualizing output of the Object Detector on the video

```

import cv2
from pprint import pprint
from matplotlib import pyplot as plt

def annotate_video(detections, input_video_path, output_video_path):
    color1=(207, 248, 64)
    color2=(255, 49, 49)
    thickness=4

    vcap = cv2.VideoCapture(input_video_path)
    width = int(vcap.get(3))
    height = int(vcap.get(4))
    fps = vcap.get(5)
    fourcc = cv2.VideoWriter_fourcc('m', 'p', '4', 'v') #codec
    video=cv2.VideoWriter(output_video_path, fourcc, fps, (width,height))

    frame_id = 0
    # Capture frame-by-frame
    # ret = 1 if the video is captured; frame is the image
    ret, frame = vcap.read()

    while ret:
        df = detections
        df = df[['fastrcnnobjectdetector.bboxes', 'fastrcnnobjectdetector.labels']][df.
↪index == frame_id]
        if df.size:
            dfLst = df.values.tolist()
            for bbox, label in zip(dfLst[0][0], dfLst[0][1]):
                x1, y1, x2, y2 = bbox
                x1, y1, x2, y2 = int(x1), int(y1), int(x2), int(y2)
                # object bbox
                img=cv2.rectangle(frame, (x1, y1), (x2, y2), color1, thickness)
                # object label
                cv2.putText(img, label, (x1, y1-10), cv2.FONT_HERSHEY_SIMPLEX, 0.9,
↪color1, thickness)
                # frame label
                cv2.putText(img, 'Frame ID: ' + str(frame_id), (700, 500), cv2.FONT_
↪HERSHEY_SIMPLEX, 1.2, color2, thickness)
                video.write(img)

            # Show every fifth frame
            if frame_id % 5 == 0:
                plt.imshow(img)
                plt.show()

```

(continues on next page)

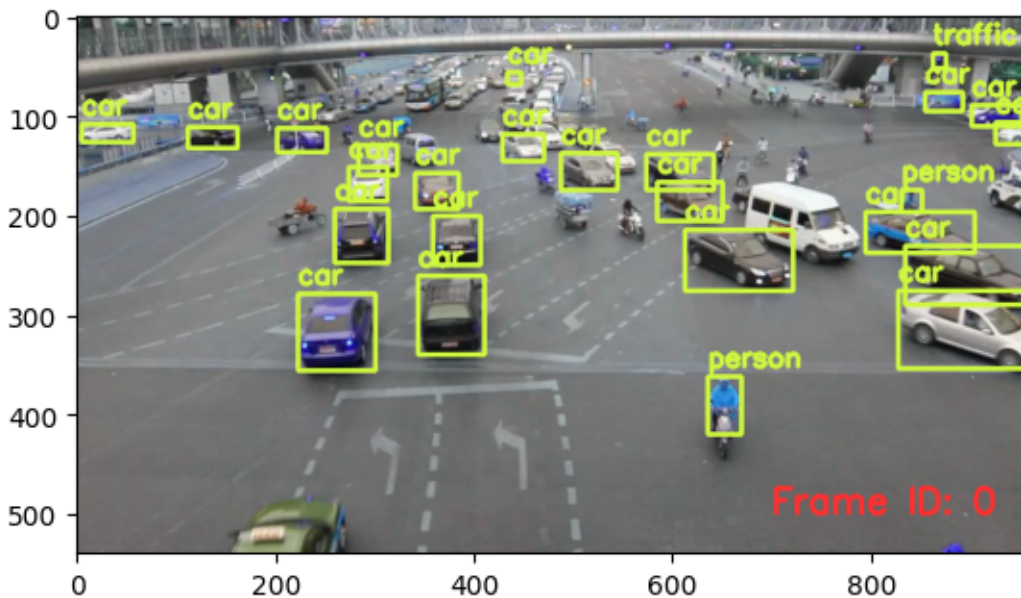
(continued from previous page)


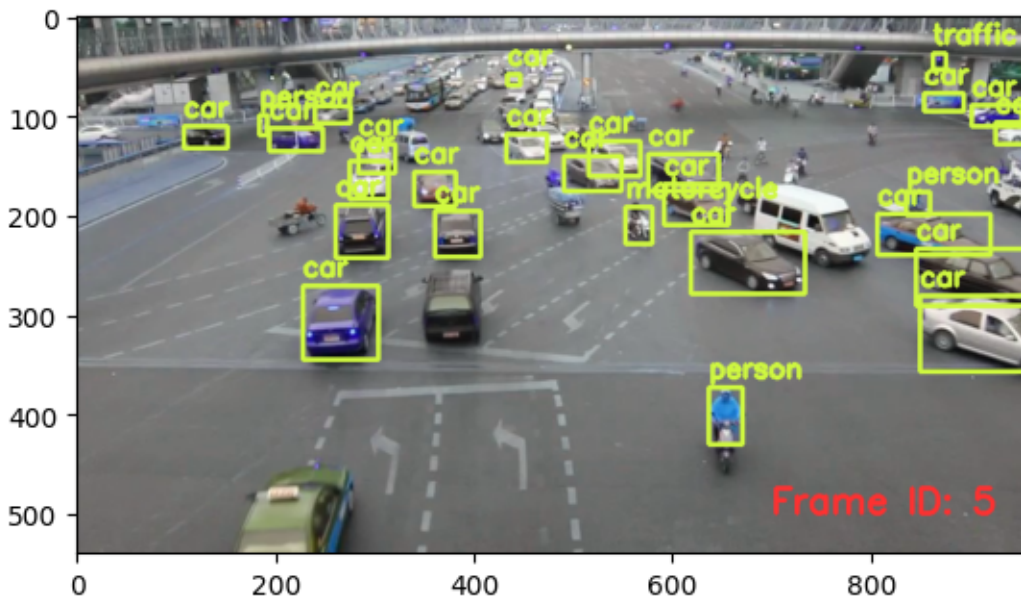
```
frame_id+=1
ret, frame = vcap.read()


video.release()
vcap.release()
```

```
from ipywidgets import Video, Image
input_path = 'ua_detrac.mp4'
output_path = 'video.mp4'

dataframe = response.batch.frames
annotate_video(dataframe, input_path, output_path)
Video.from_file(output_path)
```







```
Video(value=b'\x00\x00\x00\x1cftypisom\x00\x00\x02\x00isomiso2mp41\x00\x00\x00\x08free\x00\n\x85\xfe...')
```

11.8 Dropping an User-Defined Function (UDF)

```
cursor.execute("DROP UDF FastRCNNObjectDetector;")
response = cursor.fetch_all()
print(response)
```

```
@status: ResponseStatus.SUCCESS
@batch:
0
0 UDF FastRCNNObjectDetector successfully dropped
@query_time: 0.030028007924556732
```


MOVIE ANALYSIS

12.1 Start EVA Server

We are reusing the start server notebook for launching the EVA server

```
!wget -nc "https://raw.githubusercontent.com/georgia-tech-db/eva/master/tutorials/00-  
↪start-eva-server.ipynb"  
%run 00-start-eva-server.ipynb  
cursor = connect_to_server()
```

```
File '00-start-eva-server.ipynb' already there; not retrieving.
```

```
[ -z "$(lsof -ti:5432)" ] || kill -9 "$(lsof -ti:5432)"
```

```
nohup eva_server > eva.log 2>&1 &
```

```
Note: you may need to restart the kernel to use updated packages.
```

12.2 Video Files

getting some video files to test

```
# A video of a happy person  
!wget -nc https://www.dropbox.com/s/gzfhwmb7u804zy/defhappy.mp4  
  
# Adding Emotion detection  
!wget -nc https://raw.githubusercontent.com/georgia-tech-db/eva/master/eva/udfs/emotion_  
↪detector.py  
  
# Adding Face Detector  
!wget -nc https://raw.githubusercontent.com/georgia-tech-db/eva/master/eva/udfs/face_  
↪detector.py
```

```
File 'defhappy.mp4' already there; not retrieving.
```

```
File 'emotion_detector.py' already there; not retrieving.
```

```
File 'face_detector.py' already there; not retrieving.
```

12.3 Adding the video file to EVADB for analysis

```
cursor.execute('LOAD FILE "defhappy.mp4" INTO HAPPY')
response = cursor.fetch_all()
print(response)
```

```
@status: ResponseStatus.SUCCESS
@batch:
0
0 Video successfully added at location: defhappy.mp4
@query_time: 0.04395428113639355
```

12.4 Visualize Video

```
from IPython.display import Video
Video("defhappy.mp4", height=450, width=800, embed=True)
```

```
<IPython.core.display.Video object>
```

12.5 Create an user-defined function(UDF) for analyzing the frames

```
cursor.execute("""CREATE UDF IF NOT EXISTS EmotionDetector
    INPUT (frame NDARRAY UINT8(3, ANYDIM, ANYDIM))
    OUTPUT (labels NDARRAY STR(ANYDIM), scores NDARRAY FLOAT32(ANYDIM))
    TYPE Classification IMPL 'emotion_detector.py';
""")
response = cursor.fetch_all()
print(response)
cursor.execute("""CREATE UDF FaceDetector
    INPUT (frame NDARRAY UINT8(3, ANYDIM, ANYDIM))
    OUTPUT (bboxes NDARRAY FLOAT32(ANYDIM, 4),
            scores NDARRAY FLOAT32(ANYDIM))
    TYPE FaceDetection
    IMPL 'face_detector.py';
""")
response = cursor.fetch_all()
print(response)
```

```
@status: ResponseStatus.SUCCESS
@batch:
0
0 UDF EmotionDetector already exists, nothing added.
@query_time: 0.020081426948308945
```

(continues on next page)

(continued from previous page)

```
@status: ResponseStatus.FAIL
@batch:
  None
@error: UDF FaceDetector already exists.
```

12.6 Run the Face Detection UDF on video

```
cursor.execute("""SELECT id, FaceDetector(data)
                  FROM HAPPY WHERE id<10""")
response = cursor.fetch_all()
print(response)
```

```
@status: ResponseStatus.SUCCESS
@batch:
  happy.id          facedetector.bboxes facedetector.scores
0             0      [[493.4729  89.1079 769.5677 441.7005]]      [0.9997701]
1             1      [[501.57224 89.75874 773.5567 442.99564]]      [0.99984527]
2             2      [[503.33987 92.72798 773.9184 444.89667]]      [0.9998871]
3             3      [[506.64032 91.81204 774.7923 446.76813]]      [0.9994814]
4             4      [[508.46338 93.29222 777.00275 448.40146]]      [0.99958366]
5             5      [[506.592   99.09416 772.56396 445.32654]]      [0.99950814]
6             6      [[508.85898 98.99975 774.42487 450.29272]]      [0.999731]
7             7      [[512.3384  98.99459 781.2488 451.23007]]      [0.9997571]
8             8      [[513.3356  97.60012 783.34937 451.96744]]      [0.99983895]
9             9      [[514.25696 98.21023 784.0434  452.2896  ]]      [0.9998286]
10            0      [[493.4729  89.1079 769.5677 441.7005]]      [0.9997701]
11            1      [[501.57224 89.75874 773.5567 442.99564]]      [0.99984527]
12            2      [[503.33987 92.72798 773.9184 444.89667]]      [0.9998871]
13            3      [[506.64032 91.81204 774.7923 446.76813]]      [0.9994814]
14            4      [[508.46338 93.29222 777.00275 448.40146]]      [0.99958366]
15            5      [[506.592   99.09416 772.56396 445.32654]]      [0.99950814]
16            6      [[508.85898 98.99975 774.42487 450.29272]]      [0.999731]
17            7      [[512.3384  98.99459 781.2488 451.23007]]      [0.9997571]
18            8      [[513.3356  97.60012 783.34937 451.96744]]      [0.99983895]
19            9      [[514.25696 98.21023 784.0434  452.2896  ]]      [0.9998286]
@query_time: 3.7834010757505894
```

12.7 Run the Emotion Detection UDF on the outputs of the Face Detection UDF

```
cursor.execute("""SELECT id, bbox, EmotionDetector(Crop(data, bbox))
                  FROM HAPPY JOIN LATERAL UNNEST(FaceDetector(data)) AS Face(bbox,
↳conf)
                  WHERE id < 15;""")
response = cursor.fetch_all()
print(response)
```

```
@status: ResponseStatus.SUCCESS
```

```
@batch:
```

	happy.id	Face.bbox \
0	0	[493.4729, 89.1079, 769.5677, 441.7005]
1	1	[501.57224, 89.75874, 773.5567, 442.99564]
2	2	[503.33987, 92.728, 773.9184, 444.89667]
3	3	[506.64032, 91.81204, 774.7923, 446.76813]
4	4	[508.46338, 93.29221, 777.00275, 448.40146]
5	5	[506.592, 99.094154, 772.56396, 445.32654]
6	6	[508.85898, 98.99975, 774.42487, 450.29272]
7	7	[512.3384, 98.99459, 781.2488, 451.23007]
8	8	[513.3356, 97.60013, 783.34937, 451.96744]
9	9	[514.25696, 98.21022, 784.0434, 452.2896]
10	10	[515.7314, 97.13689, 786.78296, 452.3648]
11	10	[50.10141, 524.19183, 120.75999, 599.16064]
12	11	[513.3706, 96.37443, 784.8802, 452.24167]
13	12	[512.84, 95.54421, 785.3812, 453.33618]
14	13	[512.6082, 94.14093, 785.8945, 452.27637]
15	14	[512.9006, 93.912674, 786.7881, 453.10126]
16	0	[493.4729, 89.1079, 769.5677, 441.7005]
17	1	[501.57224, 89.75874, 773.5567, 442.99564]
18	2	[503.33987, 92.728, 773.9184, 444.89667]
19	3	[506.64032, 91.81204, 774.7923, 446.76813]
20	4	[508.46338, 93.29221, 777.00275, 448.40146]
21	5	[506.592, 99.094154, 772.56396, 445.32654]
22	6	[508.85898, 98.99975, 774.42487, 450.29272]
23	7	[512.3384, 98.99459, 781.2488, 451.23007]
24	8	[513.3356, 97.60013, 783.34937, 451.96744]
25	9	[514.25696, 98.21022, 784.0434, 452.2896]
26	10	[515.7314, 97.13689, 786.78296, 452.3648]
27	10	[50.10141, 524.19183, 120.75999, 599.16064]
28	11	[513.3706, 96.37443, 784.8802, 452.24167]
29	12	[512.84, 95.54421, 785.3812, 453.33618]
30	13	[512.6082, 94.14093, 785.8945, 452.27637]
31	14	[512.9006, 93.912674, 786.7881, 453.10126]

	emotiondetector.labels	emotiondetector.scores
0	happy	0.999640
1	happy	0.999674
2	happy	0.999696
3	happy	0.999686
4	happy	0.999690
5	happy	0.999709
6	happy	0.999732
7	happy	0.999721
8	happy	0.999709
9	happy	0.999718
10	happy	0.999701
11	neutral	0.998291
12	happy	0.999687
13	happy	0.999676
14	happy	0.999639
15	happy	0.999649

(continues on next page)

(continued from previous page)

```

16             happy                0.999640
17             happy                0.999674
18             happy                0.999696
19             happy                0.999686
20             happy                0.999690
21             happy                0.999709
22             happy                0.999732
23             happy                0.999721
24             happy                0.999709
25             happy                0.999718
26             happy                0.999701
27             neutral              0.998291
28             happy                0.999687
29             happy                0.999676
30             happy                0.999639
31             happy                0.999649
@query_time: 2.1181997805833817

```

```

import cv2
from pprint import pprint
from matplotlib import pyplot as plt

def annotate_video(detections, input_video_path, output_video_path):
    color1=(207, 248, 64)
    color2=(255, 49, 49)
    thickness=4

    vcap = cv2.VideoCapture(input_video_path)
    width = int(vcap.get(3))
    height = int(vcap.get(4))
    fps = vcap.get(5)
    fourcc = cv2.VideoWriter_fourcc('m', 'p', '4', 'v') #codec
    video=cv2.VideoWriter(output_video_path, fourcc, fps, (width,height))

    frame_id = 0
    # Capture frame-by-frame
    # ret = 1 if the video is captured; frame is the image
    ret, frame = vcap.read()

    while ret:
        df = detections
        df = df[['Face.bbox', 'emotiondetector.labels', 'emotiondetector.scores']][df.
→index == frame_id]
        if df.size:

            x1, y1, x2, y2 = df['Face.bbox'].values[0]
            label = df['emotiondetector.labels'].values[0]
            score = df['emotiondetector.scores'].values[0]
            x1, y1, x2, y2 = int(x1), int(y1), int(x2), int(y2)
            # object bbox
            img=cv2.rectangle(frame, (x1, y1), (x2, y2), color1, thickness)
            # object label

```

(continues on next page)

(continued from previous page)

```

        cv2.putText(img, label, (x1, y1-10), cv2.FONT_HERSHEY_SIMPLEX, 0.9, color1, ↵
        ↪thickness)
        # object score
        cv2.putText(img, str(round(score, 5)), (x1+120, y1-10), cv2.FONT_HERSHEY_
        ↪SIMPLEX, 0.9, color1, thickness)
        # frame label
        cv2.putText(img, 'Frame ID: ' + str(frame_id), (700, 500), cv2.FONT_HERSHEY_
        ↪SIMPLEX, 1.2, color2, thickness)
        video.write(img)

        # Show every fifth frame
        if frame_id % 5 == 0:
            plt.imshow(img)
            plt.show()

        frame_id+=1
        ret, frame = vcap.read()

    video.release()
    vcap.release()


```



```



from ipywidgets import Video, Image
input_path = 'defhappy.mp4'
output_path = 'video.mp4'



dataframe = response.batch.frames
annotate_video(dataframe, input_path, output_path)

```









CUSTOM MODEL TUTORIAL

13.1 Start EVA server

We are reusing the start server notebook for launching the EVA server.

```
!wget -nc "https://raw.githubusercontent.com/georgia-tech-db/eva/master/tutorials/00-  
↪start-eva-server.ipynb"  
%run 00-start-eva-server.ipynb  
cursor = connect_to_server()
```

```
File '00-start-eva-server.ipynb' already there; not retrieving.
```

```
[ -z "$(lsof -ti:5432)" ] || kill -9 "$(lsof -ti:5432)"
```

```
nohup eva_server > eva.log 2>&1 &
```

```
Note: you may need to restart the kernel to use updated packages.
```

13.2 Download custom user-defined function (UDF), model, and video

```
# Download UDF  
!wget -nc https://www.dropbox.com/s/lharq14izp08bfz/gender.py  
  
# Download built-in Face Detector  
!wget -nc https://raw.githubusercontent.com/georgia-tech-db/eva/master/eva/udfs/face_  
↪detector.py  
  
# Download models  
!wget -nc https://www.dropbox.com/s/0y291evpqdfmv2z/gender.pth  
  
# Download videos  
!wget -nc https://www.dropbox.com/s/f5447euuuis1vdy/short.mp4
```

```
File 'gender.py' already there; not retrieving.
```

```
File 'face_detector.py' already there; not retrieving.
```

```
File 'gender.pth' already there; not retrieving.
```

```
File 'short.mp4' already there; not retrieving.
```

13.3 Load video for analysis

```
cursor.execute("DROP TABLE TIKTOK;")
response = cursor.fetch_all()
print(response)
cursor.execute("LOAD FILE 'short.mp4' INTO TIKTOK;")
response = cursor.fetch_all()
print(response)
cursor.execute("""SELECT id FROM TIKTOK WHERE id < 5""")
response = cursor.fetch_all()
print(response)
```

```
@status: ResponseStatus.SUCCESS
@batch:
0
0 Table Successfully dropped: TIKTOK
@query_time: 0.036311471834778786
```

```
@status: ResponseStatus.SUCCESS
@batch:
0
0 Video successfully added at location: short.mp4
@query_time: 0.046881699934601784
```

```
@status: ResponseStatus.SUCCESS
@batch:
tiktok.id
0 0
1 1
2 2
3 3
4 4
@query_time: 0.2178105916827917
```

13.4 Visualize Video

```
from IPython.display import Video
Video("short.mp4", embed=True)
```

```
<IPython.core.display.Video object>
```

13.5 Create GenderCNN and FaceDetector UDFs

```

cursor.execute("""DROP UDF GenderCNN;""")
response = cursor.fetch_all()
print(response)

cursor.execute("""CREATE UDF IF NOT EXISTS
                  GenderCNN
                  INPUT (data NDARRAY UINT8(3, 224, 224))
                  OUTPUT (label TEXT(10))
                  TYPE Classification
                  IMPL 'gender.py';
                """)
response = cursor.fetch_all()
print(response)

cursor.execute("""CREATE UDF IF NOT EXISTS
                  FaceDetector
                  INPUT (frame NDARRAY UINT8(3, ANYDIM, ANYDIM))
                  OUTPUT (bboxes NDARRAY FLOAT32(ANYDIM, 4),
                          scores NDARRAY FLOAT32(ANYDIM))
                  TYPE FaceDetection
                  IMPL 'face_detector.py';
                """)
response = cursor.fetch_all()
print(response)

```

```

@status: ResponseStatus.SUCCESS
@batch:
0
0 UDF GenderCNN successfully dropped
@query_time: 0.02437913604080677

```

```

@status: ResponseStatus.SUCCESS
@batch:
0
0 UDF GenderCNN successfully added to the database.
@query_time: 2.822970863431692
@status: ResponseStatus.SUCCESS
@batch:
0
0 UDF FaceDetector already exists, nothing added.
@query_time: 0.018602540716528893

```


13.6 Run Face Detector on video

```
cursor.execute("""SELECT id, FaceDetector(data).bboxes
                FROM TIKTOK WHERE id < 10""")
response = cursor.fetch_all()
print(response)
```

```
@status: ResponseStatus.SUCCESS
@batch:
    tiktok.id          facedetector.bboxes
0           0      [[ 90.70622 208.44966 281.64642 457.68872]]
1           1      [[ 91.01816 208.27583 281.0808 457.91995]]
2           2      [[ 90.358536 207.3743 283.4399 457.96234 ]]
3           3      [[ 90.694214 207.56027 284.37817 458.6282  ]]
4           4      [[ 90.684944 208.98653 282.1281 460.90894 ]]
5           5      [[ 89.47423 209.38083 283.45938 460.58548]]
6           6      [[ 88.50081 208.31546 283.29172 461.8374  ]]
7           7      [[ 89.838646 206.07619 282.93942 464.7494  ]]
8           8      [[ 90.18522 224.35588 281.29733 469.89603]]
9           9      [[ 94.34447 234.13255 279.6476 468.85303]]
@query_time: 1.2682587169110775
```

13.7 Composing UDFs in a query

Detect gender of the faces detected in the video by composing a set of UDFs (GenderCNN, FaceDetector, and Crop)

```
cursor.execute("""SELECT id, bbox, GenderCNN(Crop(data, bbox))
                FROM TIKTOK JOIN LATERAL UNNEST(FaceDetector(data)) AS Face(bbox,
↳conf)
                WHERE id < 10;""")
response = cursor.fetch_all()
print(response)
```

```
@status: ResponseStatus.SUCCESS
@batch:
    tiktok.id          Face.bbox  gendercnn.label
0           0      [90.70622, 208.44966, 281.64642, 457.68872]      female
1           1      [91.01816, 208.27583, 281.0808, 457.91995]      female
2           2      [90.358536, 207.3743, 283.4399, 457.96234]      female
3           3      [90.694214, 207.56027, 284.37817, 458.6282]      female
4           4      [90.684944, 208.98653, 282.1281, 460.90894]      female
5           5      [89.47423, 209.38083, 283.45938, 460.58548]      female
6           6      [88.50081, 208.31546, 283.29172, 461.8374]      female
7           7      [89.838646, 206.07619, 282.93942, 464.7494]      female
8           8      [90.18522, 224.35588, 281.29733, 469.89603]      female
9           9      [94.34447, 234.13255, 279.6476, 468.85303]      female
@query_time: 0.8337160255759954
```


13.8 Visualize Output

```

import cv2
from matplotlib import pyplot as plt

def annotate_video(detections, input_video_path, output_video_path):
    color=(207, 248, 64)
    thickness=4

    vcap = cv2.VideoCapture(input_video_path)
    width = int(vcap.get(3))
    height = int(vcap.get(4))
    fps = vcap.get(5)
    fourcc = cv2.VideoWriter_fourcc(*'MP4V') #codec
    video=cv2.VideoWriter(output_video_path, fourcc, fps, (width,height))

    frame_id = 0
    # Capture frame-by-frame
    ret, frame = vcap.read() # ret = 1 if the video is captured; frame is the image

    while ret:
        df = detections
        df = df[['Face.bbox', 'gendercnn.label']][df['tiktok.id'] == frame_id]

        if df.size:
            for bbox, label in df.values:
                x1, y1, x2, y2 = bbox
                x1, y1, x2, y2 = int(x1), int(y1), int(x2), int(y2)
                img=cv2.rectangle(frame, (x1, y1), (x2, y2), color, thickness) # object_
                cv2.putText(img, str(label), (x1, y1-10), cv2.FONT_HERSHEY_SIMPLEX, 0.9,
                color, thickness-1) # object label
                video.write(img)

            # Show every fifth frame
            if frame_id % 5 == 0:
                plt.imshow(img)
                plt.show()

            if frame_id == 20:
                return

        frame_id+=1
        ret, frame = vcap.read()

    video.release()
    vcap.release()

```

```


#!/pip install ipywidgets
from ipywidgets import Video
input_path = 'short.mp4'
output_path = 'annotated_short.mp4'



```

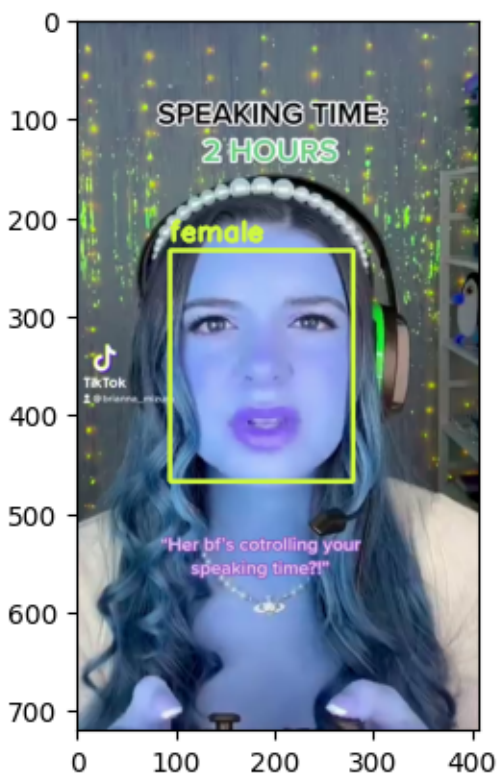

(continues on next page)

(continued from previous page)

```
dataframe = response.batch.frames  
annotate_video(dataframe, input_path, output_path)
```







EVA QUERY LANGUAGE REFERENCE

EVA Query Language (EVAQL) is derived from SQL. It is tailored for video analytics. EVAQL allows users to invoke deep learning models in the form of user-defined functions (UDFs).

Here is an example where we first define a UDF wrapping around the FastRCNN object detection model. We then issue a query with this function to detect objects.

```
--- Create an user-defined function wrapping around FastRCNN ObjectDetector
CREATE UDF IF NOT EXISTS FastRCNNObjectDetector
INPUT (frame NDARRAY UINT8(3, ANYDIM, ANYDIM))
OUTPUT (labels NDARRAY STR(ANYDIM), bboxes NDARRAY FLOAT32(ANYDIM, 4),
        scores NDARRAY FLOAT32(ANYDIM))
TYPE Classification
IMPL 'eva/udfs/fastrcnn_object_detector.py';

--- Use the function to retrieve frames that contain more than 3 cars
SELECT id FROM MyVideo
WHERE Array_Count(FastRCNNObjectDetector(data).label, 'car') > 3
ORDER BY id;
```

This page presents a list of all the EVAQL statements that you can leverage in your Jupyter Notebooks.

14.1 LOAD

14.1.1 LOAD VIDEO

```
LOAD FILE 'test_video.mp4' INTO MyVideo;

--- Alternate syntax that explicitly specifies format
LOAD FILE 'dummy.avi' INTO MyVideo WITH FORMAT VIDEO;
```

- **test_video.mp4** is the location of the video file in the filesystem on the client.
- **MyVideo** is the name of the table in EVA where this video is loaded. Subsequent queries over the video must refer to this table name.

When a video is loaded, there is no need to specify the schema for the video table. EVA automatically generates the following schema with two columns: `id` and `data`, that correspond to the frame id and frame content (in Numpy format).

14.1.2 LOAD CSV

To **LOAD** a CSV file, we need to first specify the table schema.

```
CREATE TABLE IF NOT EXISTS MyCSV (  
    id INTEGER UNIQUE,  
    frame_id INTEGER,  
    video_id INTEGER,  
    dataset_name TEXT(30),  
    label TEXT(30),  
    bbox NDARRAY FLOAT32(4),  
    object_id INTEGER  
);  
  
LOAD FILE 'test_metadata.csv' INTO MyCSV WITH FORMAT CSV;
```

- **test_metadata.csv** needs to be loaded onto the server using **LOAD** statement.
- The CSV file may contain additional columns. EVA will only load the columns listed in the defined schema.
- **WITH FORMAT CSV** is required to distinguish between videos and CSV files.

14.2 SELECT

14.2.1 SELECT FRAMES WITH PREDICATES

Search for frames with a car

```
SELECT id, frame  
FROM MyVideo  
WHERE ['car'] <@ FastRCNNObjectDetector(frame).labels  
ORDER BY id;
```

Search frames with a pedestrian and a car

```
SELECT id, frame  
FROM MyVideo  
WHERE ['pedestrian', 'car'] <@ FastRCNNObjectDetector(frame).labels;
```

Search for frames containing greater than 3 cars

```
SELECT id FROM MyVideo  
WHERE Array_Count(FastRCNNObjectDetector(data).label, 'car') > 3  
ORDER BY id;
```

14.2.2 SELECT WITH MULTIPLE UDFS

Compose multiple user-defined functions in a single query to construct semantically complex queries.

```
SELECT id, bbox, EmotionDetector(Crop(data, bbox))
FROM HAPPY JOIN LATERAL UNNEST(FaceDetector(data)) AS Face(bbox, conf)
WHERE id < 15;
```

14.3 EXPLAIN

14.3.1 EXPLAIN QUERY

List the query plan associated with a EVAQL query

Append EXPLAIN in front of the query to retrieve the plan.

```
EXPLAIN SELECT CLASS FROM TAIPAI;
```

14.4 SHOW

14.4.1 SHOW UDFS

List the registered user-defined functions

```
SHOW UDFS;
```

14.5 CREATE

14.5.1 CREATE TABLE

To create a table, specify the schema of the table.

```
CREATE TABLE IF NOT EXISTS MyCSV (
    id INTEGER UNIQUE,
    frame_id INTEGER,
    video_id INTEGER,
    dataset_name TEXT(30),
    label TEXT(30),
    bbox NDARRAY FLOAT32(4),
    object_id INTEGER
);
```

14.5.2 CREATE UDF

To register an user-defined function, specify the implementation details of the UDF.

```
CREATE UDF IF NOT EXISTS FastRCNNObjectDetector
INPUT (frame NDARRAY UINT8(3, ANYDIM, ANYDIM))
OUTPUT (labels NDARRAY STR(ANYDIM), bboxes NDARRAY FLOAT32(ANYDIM, 4),
        scores NDARRAY FLOAT32(ANYDIM))
TYPE Classification
IMPL 'eva/udfs/fastrcnn_object_detector.py';
```

14.5.3 CREATE MATERIALIZED VIEW

To create a view with materialized results – like the outputs of deep learning model, use the following template:

```
CREATE MATERIALIZED VIEW UAETRAC_FastRCNN (id, labels) AS
SELECT id, FastRCNNObjectDetector(frame).labels
FROM UAETRAC
WHERE id<5;
```

14.6 DROP

14.6.1 DROP TABLE

```
DROP TABLE DETRACVideo;
```

14.6.2 DROP UDF

```
DROP UDF FastRCNNObjectDetector;
```

14.7 INSERT

14.7.1 INSERT INTO TABLE

Insert a tuple into a table.

```
INSERT INTO MyVideo (id, data) VALUES
(1,
 [[40, 40, 40] , [40, 40, 40]],
 [[40, 40, 40] , [40, 40, 40]]);
```


14.8 RENAME

14.8.1 RENAME TABLE

```
RENAME TABLE MyVideo TO MyVideo1;
```


USER-DEFINED FUNCTIONS

This section provides an overview of how you can create and use a custom user-defined function (UDF) in your queries. For example, you could write an UDF that wraps around a PyTorch model.

15.1 Part 1: Writing a custom UDF

Illustrative UDF implementation

During each step, use [this UDF implementation](#) as a reference.

1. Create a new file under *udfs/* folder and give it a descriptive name. eg: *fastrcnn_object_detector.py*, *midas_depth_estimator.py*.

Note: UDFs packaged along with EVA are located inside the *udfs* folder.

2. Create a Python class that inherits from *PytorchClassifierAbstractUDF*.
 - The *PytorchClassifierAbstractUDF* is a parent class that defines and implements standard methods for model inference.
 - *_get_predictions()* - an abstract method that needs to be implemented in your child class.
 - *classify()* - A method that receives the frames and calls the *_get_predictions()* implemented in your child class. Based on GPU batch size, it also decides whether to split frames into chunks and performs the accumulation.
 - Additionally, it contains methods that help in:
 - Moving tensors to GPU
 - Converting tensors to numpy arrays.
 - Defining the *forward()* function that gets called when the model is invoked.
 - Basic transformations.

You can however **choose to override** these methods depending on your requirements.

3. A typical UDF class has the following components:
 - *__init__()* constructor:
 - Define variables here that will be required across all methods.
 - Model loading happens here. You can choose to load custom models or models from torch.

* **Example of loading a custom model:**

```
custom_model_path = os.path.join(EVA_DIR, "data", "models", "vehicle_
↳make_predictor", "car_recognition.pt")
self.car_make_model = CarRecognitionModel()
self.car_make_model.load_state_dict(torch.load(custom_model_path))
self.car_make_model.eval()
```

* **Example of loading a torch model:**

```
self.model = torchvision.models.detection.fasterrcnn_resnet50_
↳fpn(pretrained=True)
self.model.eval()
```

• *labels()* method:

- This should return a list of strings that your model aims to target.
- The index of the list is the value predicted by your model.

• *_get_predictions()* method:

- This is where all your model inference logic goes.
- While doing the computations, keep in mind that each call of this method is with a batch of frames.
- **Output from each invoke of the model needs to be appended to a dataframe and returned as follows:**

```
predictions = self.model(frames)
outcome = pd.DataFrame()
for prediction in predictions:

    ## YOUR INFERENCE LOGIC

    # column names depend on your implementation
    outcome = outcome.append(
        {
            "labels": pred_class,
            "scores": pred_score,
            "boxes": pred_boxes
        },
        ignore_index=True)
```

In case you have any other functional requirements (defining custom transformations etc.) you can choose to add more methods. Make sure each method you write is clear, concise and well-documented.

15.2 Part 2: Registering and using the UDF in queries

Now that you have implemented your UDF we need to register it in EVA. You can then use the function in any query.

1. Register the UDF with a query that follows this template:

```
CREATE UDF [ IF NOT EXISTS ] <name> INPUT ( [ <arg_name> <arg_data_type> ] [ , ... ] ) OUTPUT ( [ <result_name> <result_data_type> ] [ , ... ] ) TYPE <udf_type_name> IMPL '<path_to_implementation>'
```

where,

- **<name>** - specifies the unique identifier for the UDF.
- **[<arg_name> <arg_data_type>] [, ...]** - specifies the name and data type of the udf input arguments. Name is kept for consistency (ignored by eva right now), arguments data type is required. ANYDIM means the shape is inferred at runtime.
- **[<result_name> <result_data_type>] [, ...]** - specifies the name and data type of the udf output arguments. Users can access a specific output of the UDF similar to access a column of a table. Eg. <name>.<result_name>
- **<udf_type_name>** - specifies the identifier for the type of the UDF. UDFs of the same type are assumed to be interchangeable. They should all have identical input and output arguments. For example, object classification can be one type.
- **<path_to_implementation>** - specifies the path to the implementation class for the UDF

Here, is an example query that registers a UDF that wraps around the 'FastRCNNObjectDetector' model that performs Object Detection.

```
CREATE UDF IF NOT EXISTS FastRCNNObjectDetector
INPUT (frame NDARRAY UINT8(3, ANYDIM, ANYDIM))
OUTPUT (labels NDARRAY STR(ANYDIM), bboxes NDARRAY FLOAT32(ANYDIM, 4),
         scores NDARRAY FLOAT32(ANYDIM))
TYPE Classification
IMPL 'eva/udfs/fastrcnn_object_detector.py';
```

- Input is a frame of type NDARRAY with shape (3, ANYDIM, ANYDIM). 3 channels and any width or height.
- **We return 3 variables for this UDF:**
 - *labels*: Predicted label
 - *bboxes*: Bounding box of this object (rectangle coordinates)
 - *scores*: Confidence scores for this prediction

A status of 0 in the response denotes the successful registration of this UDF.

3. Now you can execute your UDF on any video:

```
SELECT id, Unnest(FastRCNNObjectDetector(data)) FROM MyVideo;
```

4. You can drop the UDF when you no longer need it.

```
DROP UDF IF EXISTS FastRCNNObjectDetector;
```


CONFIGURE GPU

1. Queries in EVA use deep learning models that run much faster on a GPU as opposed to a CPU. If your workstation has a GPU, you can configure EVA to use the GPU during query execution. Use the following command to check your hardware capabilities:

```
ubuntu-drivers devices
nvidia-smi
```

A valid output from the command indicates that your GPU is configured and ready to use. If not, you will need to install the appropriate GPU driver. [This page](#) provides a step-by-step guide on installing and configuring the GPU driver in the Ubuntu Operating System.

- When installing an NVIDIA driver, ensure that the version of the GPU driver is correct to avoid compatibility issues.
 - When installing cuDNN, you will need to create an account and ensure that you get the correct *deb* files for your operating system and architecture.
2. You can run the following code in a Jupyter notebook to verify that your GPU is detected by PyTorch:

```
import torch
device = torch.device('cuda:0' if torch.cuda.is_available() else 'cpu')
print(device)
```

Output of *cuda:0* indicates the presence of a GPU. 0 indicates the index of the GPU in system. If you have multiple GPUs on your workstation, the index must be updated accordingly.

3. Now configure the executor section in *eva.yml* as follows:

```
executor:
  gpus: {'127.0.1.1': [0]}
```

Here, *127.0.1.1* is the loopback address on which the EVA server is running. 0 refers to the GPU index to be used.

CONTRIBUTING

We welcome all kinds of contributions to EVA.

- Code reviews
- Improving documentation
- Tutorials and applications
- New features

17.1 Setting up the Development Environment

First, you will need to checkout the repository from GitHub and build EVA from the source. Follow the following instructions to build EVA locally. We recommend using a virtual environment and the pip package manager.

EVA requires JAVA 11 for generating the ANTLR-based EVAQL parser.

```
git clone https://github.com/georgia-tech-db/eva.git && cd eva
python3 -m venv test_eva_db # to create a virtual environment
. test_eva_db/bin/activate
pip install --upgrade pip
sudo -E apt install -y openjdk-11-jdk openjdk-11-jre # to install JAVA
sh script/antlr4/generate_parser.sh # to generate the EVA parser
pip install -e ".[dev]"
bash script/test/test.sh # to run the test suite
```

For developers using an M1 Mac, here are some pointers for installing JAVA and to resolve multi-threading issues:

```
brew install openjdk@11 # to install openjdk 11
export JAVA_HOME="/opt/homebrew/opt/openjdk@11" # add this command in ~/.zshrc or ~/.
↪ bashrc
export OBJC_DISABLE_INITIALIZE_FORK_SAFETY=YES # to resolve multi-threading issues in
↪ macOS
```

After you have installed the package locally, you can make changes in the code base and examine their impact.

```
pip install . # reinstall EVA package to include local changes
kill -9 eva_server # kill running EVA server (if any)
eva_server& # launch EVA server with newly installed package
```

17.2 Testing

Check if your local changes broke any unit or integration tests by running the following script:

```
bash script/test/test.sh
```

If you want to run a specific test file, use the following command.

```
python -m pytest test/integration_tests/test_select_executor.py
```

Use the following command to run a specific test case within a specific test file.

```
python -m pytest test/integration_tests/test_select_executor.py -k 'test_should_load_and_
↪select_in_table'
```

17.3 Submitting a Contribution

Follow the following steps to contribute to EVA:

- Merge the most recent changes from the master branch

```
git remote add origin git@github.com:georgia-tech-db/eva.git
git pull . origin/master
```

- Run the *test script* to ensure that all the test cases pass.
- If you are adding a new EVAQL command, add an illustrative example usage in the *documentation*.
- Run the following command to ensure that code is properly formatted.

```
python script/formatting/formatter.py
```

17.4 Code Style

We use the *black* code style for formatting the Python code. For docstrings and documentation, we use *Google Pydoc format*.

```
def function_with_types_in_docstring(param1, param2) -> bool:
    """Example function with types documented in the docstring.

    Additional explanatory text can be added in paragraphs.

    Args:
        param1 (int): The first parameter.
        param2 (str): The second parameter.

    Returns:
        bool: The return value. True for success, False otherwise.
```

17.5 Troubleshooting

If the test suite fails with a *PermissionDenied* exception, update the *path_prefix* attribute under the *storage* section in the EVA configuration file (`~/ .eva/eva.yml`) to a directory where you have write privileges.

DEBUGGING

We recommend using a debugger in Visual Studio Code for debugging EVA. This tutorial presents a detailed step-by-step process of using the debugger.

18.1 Pre-requisites

Ensure that Python extensions are installed in Visual Studio Code using these [instructions](#).

18.2 Set up Debugger for EVA

When you press the debug icon, you will be given an option to create a `launch.json` file.

While creating the JSON file, you will be prompted to select the environment to be used. Select the python environment from the `command palette` at the top. If the Python environment cannot be seen in the drop-down menu, try installing the python extension, and repeat the process.

Once you select the python environment, a `launch.json` file will be created with the default configurations set to debug a simple `.py` file.

More configurations can further be added to the file, to modify the environment variables or to debug an entire folder or workspace directory. Use the following configuration in the JSON file:

```
{
  "version": "0.2.0",
  "configurations": [
    {
      "name": "Python: test_pytorch.py",
      "type": "python",
      "request": "launch",
      "program": "${workspaceFolder}/test/integration_tests/test_pytorch.py",
      "console": "integratedTerminal",
      "cwd": "${workspaceFolder}",
      "env": {"PYTHONPATH": "${workspaceRoot}"}
    }
  ]
}
```

You can modify the fields of the above JSON file as follows:

name: It is the reader-friendly name to appear in the Debug launch configuration dropdown.

type: The type of debugger to use for this launch configuration.

program: The executable or file to run when launching the debugger. In the above example, `test_integration.py` will be executed by the debugger.

env: Here you specify the environment variables. In the above example, the path for the conda environment for Eva has been specified.

Using these configuration variables, you can run the debugger both locally as well as on a remote server.

EXTENDING EVA

This document details the steps involved in adding support for a new operator (or command) in EVA. We illustrate the process using a DDL command.

19.1 Command Handler

An input query string is handled by **Parser**, **StatementToPlanConvertor**, **PlanGenerator**, and **PlanExecutor**. We discuss each part separately.

```
def execute_query(query) -> Iterator[Batch]:  
    """  
    Execute the query and return a result generator.  
    """  
    #1. parser  
    stmt = Parser().parse(query)[0]  
    #2. statement to logical plan  
    l_plan = StatementToPlanConvertor().visit(stmt)  
    #3. logical to physical plan  
    p_plan = PlanGenerator().build(l_plan)  
    #4. parser  
    return PlanExecutor(p_plan).execute_plan()
```

19.2 1. Parser

The parser firstly generate **syntax tree** from the input string, and then transform syntax tree into **statement**.

The first part of Parser is build from two ANTLR file.

19.2.1 parser/evaql

- evaql_lexer.g4 - add keywords(eg. CREATE, TABLE) under **Common Keywords**
- evaql_parser.g4
 - Add new grammar name(eg. createTable) under **ddlStatement**
 - Write a new grammar, for example:

```
createTable
: CREATE TABLE
  ifNotExists?
  tableName createDefinitions                                #columnCreateTable
;
```

- Run ``sh script/antlr4/generate_parser.sh`` after modify g4 file to generate python file for the parser.
-

The second part of parser is implemented as **parser visitor**.

19.2.2 parser/parser_visitor

- `_[cmd]_statement.py` - eg. `class CreateTable(evaql_parserVisitor)`
 - Write functions to transform each input data from syntax tree to desired type. (eg. transform Column information into a list of ColumnDefinition)
 - Write a function to construct `[cmd]Statement` and return it.
- `__init__.py` - import `_[cmd]_statement.py` and add its class to `ParserVisitor`'s parent class.

```
from src.parser.parser_visitor._create_statement import CenameTable
class ParserVisitor(CommonClauses, CreateTable, Expressions,
                    Functions, Insert, Select, TableSources,
                    Load, Upload):
```

19.2.3 parser/

- `[cmd]_statement.py` - class `[cmd]Statement`. Its constructor is called in `_[cmd]_statement.py`
- `types.py` - register new `StatementType`

19.3 2. Statement To Plan Convertor

The part transforms the statement into corresponding logical plan.

19.3.1 Optimizer

- `operators.py`
 - Define class `Logical[cmd]`, which is the logical node for the specific type of command.

```
class LogicalCreate(Operator):
    def __init__(self, video: TableRef, column_list: List[DataFrameColumn], if_not_
    exists: bool = False, children=None):
        super().__init__(OperatorType.LOGICALCREATE, children)
        self._video = video
        self._column_list = column_list
        self._if_not_exists = if_not_exists
        # ...
```


- Register new operator type to **class OperatorType**, Notice that must add it **before LOGICALDELIMITER !!!**

- `statement_to_opr_convertor.py`

- import resource

```
from src.optimizer.operators import LogicalCreate
from src.parser.rename_statement import CreateTableStatement
```

- implement `visit_[cmd]()` function, which converts statement to operator

```
# May need to convert the statement into another data type.
# The new data type is usable for excuting command.
# For example, column_list -> column_metadata_list

def visit_create(self, statement: AbstractStatement):
    video_ref = statement.table_ref
    if video_ref is None:
        LogManager().log("Missing Table Name In Create Statement",
                        LoggingLevel.ERROR)

    if_not_exists = statement.if_not_exists
    column_metadata_list = create_column_metadata(statement.column_list)

    create_opr = LogicalCreate(
        video_ref, column_metadata_list, if_not_exists)
    self._plan = create_opr
```

- modify visit function to call the right `visit_[cmd]` function

```
def visit(self, statement: AbstractStatement):
    if isinstance(statement, SelectStatement):
        self.visit_select(statement)
    #...
    elif isinstance(statement, CreateTableStatement):
        self.visit_create(statement)
    return self._plan
```

19.4 3. Plan Generator

The part transformed logical plan to physical plan. The modified files are stored under **Optimizer** and **Planner** folders.

19.4.1 planner/

- [cmd]_plan.py - class [cmd]Plan, which stored information required for rename table.

```
class CreatePlan(AbstractPlan):
    def __init__(self, video_ref: TableRef,
                  column_list: List[DataFrameColumn],
                  if_not_exists: bool = False):
        super().__init__(PlanOprType.CREATE)
        self._video_ref = video_ref
        self._column_list = column_list
        self._if_not_exists = if_not_exists
    #...
```

- types.py - register new plan operator type to PlanOprType

19.4.2 optimizer/rules

- rules.py-
 - Import operators
 - Register new rule type to **RuleType** and **Promise** (place it **before IMPLEMENTATION_DELIMITER !!**)
 - implement class Logical[cmd]ToPhysical, its member function apply() will construct a corresponding [cmd]Plan object.

```
class LogicalCreateToPhysical(Rule):
    def __init__(self):
        pattern = Pattern(OperatorType.LOGICALCREATE)
        super().__init__(RuleType.LOGICAL_CREATE_TO_PHYSICAL, pattern)

    def promise(self):
        return Promise.LOGICAL_CREATE_TO_PHYSICAL

    def check(self, before: Operator, context: OptimizerContext):
        return True

    def apply(self, before: LogicalCreate, context: OptimizerContext):
        after = CreatePlan(before.video, before.column_list, before.if_not_exists)
        return after
```

19.5 4. Plan Executor

PlanExecutor uses data stored in physical plan to run the command.

19.5.1 executor/

- [cmd]_executor.py - implement an executor that make changes in **catalog**, **metadata**, or **storage engine** to run the command.
 - May need to create helper function in CatalogManager, DatasetService, DataFrameMetadata, etc.

```
class CreateExecutor(AbstractExecutor):
    def exec(self):
        if (self.node.if_not_exists):
            # check catalog if we already have this table
            return

        table_name = self.node.video_ref.table_info.table_name
        file_url = str(generate_file_path(table_name))
        metadata = CatalogManager().create_metadata(table_name, file_url, self.node.
↪column_list)

        StorageEngine.create(table=metadata)
```

19.6 Additional Notes

Key data structures in EVA:

- **Catalog:** Records DataFrameMetadata for all tables.
 - data stored in DataFrameMetadata: name, file_url, identifier_id, schema
 - * file_url - used to access the real table in storage engine.
 - For the RENAME table command, we use the old_table_name to access the corresponding entry in metadata table, and the modified name of the table.
- **Storage Engine:**
 - API is defined in src/storage, currently only supports create, read, write.

EVA RELEASE GUIDE

20.1 Part 1: Before You Start

Make sure you have [PyPI](#) account with maintainer access to the EVA project. Create a `.pypirc` in your home directory. It should look like this:

```
[distutils]
index-servers =
  pypi
  pypitest

[pypi]
username=YOUR_USERNAME
password=YOUR_PASSWORD
```

Then run `chmod 600 ~/.pypirc` so that only you can read/write the file.

20.2 Part 2: Release Steps

1. Ensure that you're in the top-level `eva` directory.
2. Ensure that your branch is in sync with the `master` branch:

```
$ git pull origin master
```

3. Add a new entry in the Changelog for the release.

```
## [0.0.6]
### [Breaking Changes]
### [Added]
### [Changed]
### [Deprecated]
### [Removed]
```

Make sure `CHANGELOG.md` is up to date for the release: compare against PRs merged since the last release.

4. Update version to, e.g. `0.0.6` (remove the `+dev` label) in `eva/version.py`.
5. Commit these changes and create a PR:

```
git checkout -b release-v0.0.6
git add . -u
git commit -m "[RELEASE]: v0.0.6"
git push --set-upstream origin release-v0.0.6
```

6. Once the PR is approved, merge it and pull master locally.

7. Tag the release:

```
git tag -a v0.0.6 -m "v0.0.6 release"
git push origin v0.0.6
```

8. Build the source and wheel distributions:

```
rm -rf dist build # clean old builds & distributions
python3 setup.py sdist # create a source distribution
python3 setup.py bdist_wheel # create a universal wheel
```

9. Check that everything looks correct by installing the wheel locally and checking the version:

```
python3 -m venv test_evadb # create a virtualenv for testing
source test_evadb/bin/activate # activate virtualenv
python3 -m pip install dist/evadb-0.9.1-py3-none-any.whl
python3 -c "import eva; print(eva.__version__)"
```

10. Publish to PyPI

```
pip install twine # if not installed
twine upload dist/* -r pypi
```

11. A PR is automatically submitted (this will take a few hours) on [\[conda-forge/eva-feedstock\]](https://github.com/conda-forge/eva-feedstock) (<https://github.com/conda-forge/eva-feedstock>) to update the version. * A maintainer needs to accept and merge those changes.

12. Create a new release on Github. * Input the recently-created Tag Version: `v0.0.6` * Copy the release notes in `CHANGELOG.md` to the GitHub tag. * Attach the resulting binaries in `(dist/evadb-x.x.x.*)` to the release. * Publish the release.

13. Update version to, e.g. `0.9.1+dev` in `eva/version.py`.

14. Add a new changelog entry for the unreleased version in `CHANGELOG.md`:

```
## [Unreleased]
### [Breaking Changes]
### [Added]
### [Changed]
### [Deprecated]
### [Removed]
```

15. Commit these changes and create a PR:

```
git checkout -b bump-v0.9.1+dev
git add . -u
git commit -m "[BUMP]: v0.9.1+dev"
git push --set-upstream origin bump-v0.9.1+dev
```

16. Add the new tag to the EVA project on ReadTheDocs,

- Trigger a build for main to pull new tags.
- Go to the Versions tab, and Activate the new tag.
- Go to Admin/Advanced to set this tag as the new default version.
- **In Overview, make sure a build is triggered:**
 - For the tag v0.9.1
 - For latest

Credits: [Snorkel](#)

PACKAGING

This section describes practices to follow when packaging your own models or datasets to be used along with EVA.

21.1 Models

Please follow the following steps to package models:

- Create a folder with a descriptive name. This folder name will be used by the UDF that is invoking your model.
- **Place all files used by the UDF inside this folder. These are typically:**
 - Model weights (The .pt files that contain the actual weights)
 - Model architectures (The .pt files that contain model architecture information)
 - Label files (Extra files that are used in the process of model inference for outputting labels.)
 - Other config files (Any other config files required for model inference)
- Zip this folder.
- Upload the zipped folder to this [link](#) inside the models folder.

21.2 Datasets

Please follow the following steps to package datasets:

- Create a folder for your dataset and give it a descriptive name.
- This dataset folder should contain 2 sub-folders named ‘info’ and ‘videos’. For each video entry in the videos folder, there should be a corresponding CSV file in the info folder with the same name. The structure should look like:

```

anipi@blade:~/codes/eva$ tree data/datasets/bddtest/
data/datasets/bddtest/
├── info
│   ├── 00a04f658c891f94.csv
│   ├── 00a04f65af2ab984.csv
│   ├── 00a0f0083c67908e.csv
│   ├── 00a0f008a315437f.csv
│   ├── 00a2e3ca5c856cde.csv
│   ├── 00a2e3ca62992459.csv
│   ├── 00a2f5b6d4217a96.csv
│   ├── 00a395fed60c0b47.csv
│   ├── 00a820ef2b98dcf5.csv
│   └── 00a9cd6bb39be004.csv
└── videos
    ├── 00a04f658c891f94.mp4
    ├── 00a04f65af2ab984.mp4
    ├── 00a0f0083c67908e.mp4
    ├── 00a0f008a315437f.mp4
    ├── 00a2e3ca5c856cde.mp4
    ├── 00a2e3ca62992459.mp4
    ├── 00a2f5b6d4217a96.mp4
    ├── 00a395fed60c0b47.mp4
    ├── 00a820ef2b98dcf5.mp4
    └── 00a9cd6bb39be004.mp4

```

- The videos folder should contain the raw videos in a standard format like mp4 or mov.
- The info folder should contain the meta information corresponding to each video in CSV format. Each row of this CSV file should correspond to 1 unique object in a given frame. Please make sure the columns in your CSV file exactly match to these names. Here is a snapshot of a sample CSV file:

	A	B	C	D	E	F	G
1	id	frame_id	video_id	dataset_name	label	bbox	object_id
2	5699	0	3	bddtest	car	798.6408024113638, 240.2771362586605, 1052.2802666724026, 400.0394580512265	65349
3	5700	0	3	bddtest	car	739.5381062355658, 261.8937644341801, 845.1270207852194, 360	65350
4	5701	0	3	bddtest	car	707.9445727482679, 251.91685912240186, 807.7136258660508, 339.2147806004619	65351
5	5702	0	3	bddtest	car	478.717964316695, 272.0163071719262, 520.0833211152326, 298.941818301956	65352
6	5703	0	3	bddtest	car	688.5975996216332, 253.2197568376064, 731.6349495979853, 324.2814277287924	65353
7	5704	0	3	bddtest	car	645.4521625163827, 267.5229357798165, 675.64875491481, 297.71952817824376	65354

The columns represent the following:

- id - (Integer) Auto incrementing index that is unique across all files (Since the CSV files are written to the same meta table, we want it to be unique across all files)
- frame_id - (Integer) id of the frame this row corresponds to.
- video_id - (Integer) id of the video this file corresponds to.
- dataset_name - (String) Name of the dataset (should match the folder name)
- label - (String) label of the object this row corresponds to.
- bbox - (String) comma separated float values representing x1, y1, x2, y2 (top left and bottom right) coordinates of the bounding box

- object_id - (Integer) unique id for the object corresponding to this row.
- Zip this folder.
- Upload the zipped folder to this [link](#) inside the datasets folder.

Note: In the future, will provide utility scripts along with EVA to download models and datasets easily and place them in the appropriate locations.