



NIST Open Media Forensics Challenge

OpenMFC 2020 -2021 Results

Haiying Guan

Yooyoung Lee, Lukas Diduch, and Ilia Ghorbanian

Multimodal Information Group,
Information Access Division, ITL, NIST

OpenMFC2021 Workshop : Day 2, Wednesday, Dec. 8, 2021



Disclaimer

- Certain commercial equipment, instruments, software, or materials are identified in this article in order to specify the experimental procedure adequately. Such identification is not intended to imply recommendation or endorsement by NIST, nor is it intended to imply that the equipment, instruments, software or materials are necessarily the best available for the purpose.
- The views, opinions and/or findings expressed are those of the author and should not be interpreted as representing the official views or policies of the Department of Defense or the U.S. Government.
- All images, graphs, and charts are original works created for DARPA MediFor as well as other NIST Programs.

Acknowledgement

- NIST contributors
 - Jonathan Fiscus
 - Timothee Kheyrkhah
 - Peter Fontana
 - Jesse G. Zhang
- External collaborators:
 - Prof. Siwei Lyu in University at Buffalo
 - Prof. Jennifer Newman, Prof. Yong Guan, Li Lin from Iowa State University, and Prof. Roy A. Maxion from Carnegie Mellon University

Outline

- OpenMFC 2020-2021 Evaluation Dataset
- OpenMFC 2020-2021 Result and analysis
 - IMDL
 - VMD
 - IGMD
 - VGMD

OpenMFC 2020-2021 Evaluation Tasks

- Image Manipulation Detection and Localization (IMDL)
 - To detect if the image has been manipulated, and then to spatially localize the manipulated region
- Video Manipulation Detection (VMD)
 - To detect if the video has been manipulated
- Image GAN Manipulation Detection (IGMD)
 - To detect GAN-manipulated images (e.g., created by a GAN model, locally/globally modified by a GAN filter/operation, etc.)
- Video GAN/Deepfakes Manipulation Detection (VGMD)
 - To detect GAN/Deepfakes manipulated videos

Six OpenMFC 2020-2021 Leaderboards

<https://mfc.nist.gov/#pills-leaderboard>

- Image Manipulation Detection and Localization:
 - Image Only (IMDL-IO)
 - Image + Metadata (IMDL-IM)
- Video Manipulation Detection:
 - Video Only (VMD-VO)
 - Video + Metadata (VMD-VM)
- GAN Manipulation Detection:
 - Image Only (IGMD)
 - Video Only (VGMD)

OpenMFC 2020-2021

Evaluation Datasets

- OpenMFC 2020-2021 IMDL: MFC19 EP1 Image
- OpenMFC 2020-2021 VMD: MFC19 EP1 Video
- OpenMFC 2020-2021 IGMD: MFC18 GAN Challenge Image
- OpenMFC 2020-2021 VGMD: MFC18 GAN Challenge Video

Concepts in the Data Annotation


- Journal
 - Directed graph
- Node
- Link
- High Provenance (HP)
 - Original image
- Manipulated image
- Base Image – start node
- Final Image – end node
- Donor Image – donor node
- Probe – test sample



Manipulated image



Original image/Base Image



An Image Journal example



(a) Base Image




(b) Donor Image 1



(c) Donor Image 2





(d) Final manipulated Image






User Guide for NIST Media Forensic Challenge (MFC) Datasets, Guan, Haiying; Delgado, Andrew P.; Lee, Yooyoung; Yates, Amy N.; Zhou, Daniel F., Kheyrkhah, Timothee; Fiscus, Jonathan G., NIST Interagency/Internal Report (NISTIR) Number 8377, July 2021, Available at <https://doi.org/10.6028/NIST.IR.8377>.



Cross-Year comparisons: MFC Evaluation Dataset History



**Kick-off
2016 Dataset**





**Provenance
Auto Journaling Tool (JT)
Nimble Challenge 2017**





**New Manipulations
(CGI, Recapture, ...)
Extended JT, AutoJT**

MFC 2018



- Camera ID Eval. datasets
- Video Temporal Spatial
- Additional Manipulation Operations (GAN etc.)
- Extended JT, AutoJT

MFC 2019



Special study data

- Compression
- Global Blur
- Single Operation
- Social Media Laundering
- Frame Drop/Dup.

MFC 2020

OpenMFC IMDL Evaluation Dataset Property (1)

- MFC19 Image EP1: 75 manipulation operations
- Unique journal link count – no duplication count using link




Figure: OpenMFC 2020-2021 image eval. operation histogram of unique journal link (partial)

OpenMFC IMDL Evaluation Dataset Property (2)

- MFC19 Image EP1: 75 manipulation operations
- Test case (probe) count – has duplication counts using probe




Figure: OpenMFC 2020-2021 image eval. operation histogram of probe count (partial)

OpenMFC 2020-2021 Evaluation Dataset Summary

Table: NIST released OpenMFC datasets

Task(s)	NC16 Kickoff	NC17 EP1	MFC18 EP1	MFC19 EP1	MFC20 EP1
Image	1.1K	4K	17K	16K	20K
Video	-	0.36K	1K	1.5K	2.5K
GAN Image	-	-	1.3K	-	-
GAN Video	-	-	118	-	-

Red Color: indicates the datasets used for in the OpenMFC 2020-2021 evaluation.


Task:

Image Manipulation Detection and Localization (IMDL)

- IMDL: to detect if the image has been manipulated, and then to spatially localize the manipulated region
- Conditions:
 - Image Only (IO)
 - Image and Metadata (IM)
- Outline:
 - MediFor MFC results
 - OpenMFC participant submission file
 - OpenMFC results
 - Analysis

Detection System Evaluation Metrics


- Evaluate the accuracy of a system output (e.g., confidence score)
- Evaluation metrics
 - ROC (Receiver Operating Characteristic)
 - AUC (Area Under Curve)
 - CD (Correct Detection) @ FAR = 0.05



MFC19 IMD-IO Results

Figure: IMD on MFC19 EP1, full data


- OpenMFC 2020-2021 dataset
 - MFC19 EP1 dataset
 - 16K probe images
- 47 analytic systems from 13 teams
 - Condition: Image Only (IO)
 - Highest AUC: 0.866;
CD@FAR=0.05: 0.456



MFC Cross Year IMD-IO Performance Comparison

- Not apples to apples comparison: dataset difficulty increases each year

Image Manipulation Detection Performance




NIST OpenMFC Submission File: IMDL

- Task: IMDL
- Submission file name format:
 - anyfilename.tar.gz
- e.g., OpenMFC20_Image_DJpegDetection_TeamID_Date.tar.gz
 - 16030 rows
 - Example: the first several rows of the submission file
 - ProbeFileID|ConfidenceScore|OutputProbeMaskFileName|ProbeStatus|ProbeOptOutPixelValue
 - 00003e6a1efc7022da825396dc680343|0.29023||Processed|
 - 000d547cd11a767dab2b9bf56991bb0c|0.332853||Processed|
 - 0010d95b9edb40516bf13973cc2b9dbf|0.776651||Processed|
 - 001631ab2323cda12f9eae60c8daf2a8|0.221767||Processed|
 -

NIST OpenMFC Leaderboard: IMDL-IO

<https://mfc.nist.gov/#pills-leaderboard>

RANK	SUBMISSION ID	SUBMISSION DATE	TEAM NAME	SYSTEM NAME	AUC	CDR@0.05FAR	AVERAGE OPTIMAL MCC
1	63	2021-06-07 11:26:58	Mayachitra	test1june6	0.993707	0.972	
2	10	2020-11-05 21:53:02	UIIA	naive-efficient	0.616186	0.071351	
3	67	2021-06-08 00:51:16	UIIA	testIMDL	0.5	0.05	
4	81	2021-06-26 00:31:16	UIIA	testIMDL			0.0553688699305928



- For the IMDL-IO detection task, the best public team's AUC is 0.99 compared with 0.87 in MFC19.
- For the IMDL-IO localization task, the best public team's MCC score is 0.055 compared with 0.224 in MFC19.

OpenMFC Performance: Testing vs. Training




Figure: Team1 system performance without using the MFC19 dataset as training.




Figure : Team1 system performance using the MFC19 dataset as training.

OpenMFC Blackbox Evaluation Performance: Public Team




Figure: Team2 system performance without using MFC19 dataset as training.

Task:


Video Manipulation Detection (VMD)

- VMD: to detect if the video has been manipulated
- Conditions:
 - Video Only (VO)
 - Video and Metadata (VM)
- Outline:
 - MediFor MFC results
 - OpenMFC Participant submission file
 - OpenMFC results

MediFor MFC19 Video Manipulation Detection Results


- OpenMFC 2020-2021 dataset
 - MFC19 EP1 dataset
 - 1.5K probe videos
- 19 analytic systems from 8 teams
 - Condition:
 - Video Only (VO)
 - Video + metadata (VM)
 - Highest AUC: 0.999;
CD@FAR =0.05: 0.997

Figure: VMD on MFC19 EP1, full data



MediFor MFC Year-to-year VMD Performance Comparison

- Not apples to apples comparison: dataset complexity also increases each year



OpenMFC Submission Files: VMD

- Task: VMD
- Submission file name : e.g., OpenMFC20_Video_SystemID_TeamID_Date.tar.gz
 - 1531 rows
 - Example: the first several rows of the submission file
 - ProbeFileID|ConfidenceScore|ProbeStatus
 - 001bd1016363c47079a6165535ae7145|0.136152726|Processed
 - 00a68871d70ada52a8ef2be33df178b9|0.633048435|Processed
 - 0109fcabf33f12ab5e947afb3950c602|0.535061206|Processed
 -

Task:

Image GAN Manipulation Detection (IGMD)

- IGMD: To detect GAN-manipulated images (e.g., created by a GAN model, locally/globally modified by a GAN filter/operation, etc.).
- Conditions:
 - Image Only (IO)
- Outline:
 - IGMD evaluation Dataset
 - MediFor MFC results
 - OpenMFC Participant submission file
 - OpenMFC results

IGMD Datasets


- Manipulations in the MFC18 Image GAN evaluation dataset:
 - Face swap – GAN vs. real face
 - GAN Fill, GAN Erase
 - CFA camera model
- Reference definition
 - Manipulated by GAN operation (with/without other operations) – IsTarget = ‘Y’
 - Manipulated by other operations (not GAN) – IsTarget = ‘Y’
 - NotManipulated – IsTarget = ‘N’

NIST Data Sets	Major operations	Journals	Probe	Date
MFC18 Eval GAN Image Full (total 1340 probes)	CFA Camera Model	89	191	06/30/2018
	ErasureByGAN	86	200	
	Face Swap	37	640	
	GAN Fill	55	110	
	World Image	-	199	

GAN Full Image Dataset: A face swap journal example


- The journal was created by PAR Government Systems Corporation

11 donor
33 manipulated



MediFor MFC18 IGMD-IO ROC

Full Scoring (25)



MediFor MFC18 IGMD-IO AUC




Figure: MediFor GAN Image AUC: Holistic vs. Opt-In

* Distinct bars may represent distinct systems.

MediFor MFC18 IGMD-IO CD@FAR=0.05




Figure: MediFor GAN Image CD@FAR=0.05: Holistic vs. Opt-In

- The blue bars are the evaluation results of the default systems on the full testing set.
- The orange bars are the evaluation results of the OptIn systems which selected a subset of testing samples for the evaluation only.

OpenMFC Submission Files: IGMD

- Task: IGMD
- Submission file name : e.g., OpenMFC20_ImageGAN_SystemID_TeamID_Date.tar.gz
 - 1341 rows
 - Example: the first several rows of the submission file
 - ProbeFileID | ConfidenceScore | ProbeStatus
 - 002d809170f2b761ea8136369bddbb1c | 0.360968997 | Processed
 - 002e5f4b081c3b41d93e1414f40a588d | 0.354781056 | Processed
 - 004b354f88f064e2ac9ff555abf3ee94 | 0.11301657 | Processed
 -

OpenMFC Leaderboard: IGMD-IO

<https://mfc.nist.gov/#pills-leaderboard>

RANK	SUBMISSION ID	SUBMISSION DATE	TEAM NAME	SYSTEM NAME	AUC	CDR@0.05FAR
1	90	2021-07-10 09:56:14	UIIA	test	0.689716	0.207018
2	93	2021-07-30 18:13:11	UIIA	test	0.683956	0.187135
3	75	2021-06-14 04:12:30	UBMDFL_IGMD	dry-run	0.554261	0.012865
4	52	2021-06-01 15:41:20	UBMDFL_IGMD	dry-run	0.547125	0.009357
5	82	2021-06-23 09:21:26	UIIA	ptchatt	0.500033	0.051077




Figure : System performance on the MFC18 Image GAN dataset

- For the IGMD task, the best AUC score is 0.69, compared with 0.79 in MFC18.

Task:

Video GAN Manipulation Detection (VGMD)

- VGMD: To detect GAN/Deepfakes manipulated videos specifically while not detecting other forms of manipulations
- Conditions:
 - Video Only (VO)
- Outline:
 - VGMD dataset
 - Participant submission file
 - MediFor MFC results
 - OpenMFC results

VGMD Datasets

- Manipulations in the MFC18 video GAN evaluation dataset:
 - Deepfakes
 - Frame drop
 - GAN Inpainting
- Reference definition
 - Manipulated by GAN operation (with/without other operations) – IsTarget = ‘Y’
 - Manipulated by other operations (not GAN) – IsTarget = ‘Y’
 - NotManipulated – IsTarget = ‘N’

NIST Data Sets	Major operations	Journals	Probe	Date
MFC18 Eval GAN Video (total 118 probes)	PAR Deepfake	-	18	06/30/2018
	Dever Deepfake	-	40	
	Michigan GAN dropframe	-	20	
	PAR GAN Inpainting	-	20	
	World video	-	20	

GAN Video Example (1)

- The videos was collected and manipulated by UC Denver's team

Manipulated

85edfcdfed8e2101a6a78936bc4be733.mp4


04572fb8d0156500fb82415ce1ed55aa.mp4

OpenMFC Submission Files: VGMD

- Task: VGMD
- Submission file name : e.g., OpenMFC20_VideoGAN_SystemID_TeamID_Date.tar.gz
 - 119 rows
 - Example: the first several rows of the submission file
 - ProbeFileID|ConfidenceScore|ProbeStatus
 - 00b85e9c14b8c6d02e57255fc71c84c0|0.752036187|Processed
 - 01a61a91f9cdf22c62fa29c299d30778|0.287194415|Processed
 - 0508c1016780c9aeb4b4aac1b1bac5e9|0.328568837|Processed
 - 0744b28204e95f9a36dadeceb36a576c2|0.433211283|Processed
 - 08a6f28deff8d5fa885075b852e9f08f|0.23771142|Processed
 -

MediFor MFC18 VGMD-IO ROC

Full Scoring (7)



MediFor MFC18 VGMD-IO AUC




Figure: MediFor Video GAN AUC

5 teams: Team 1 and 3 only have default systems (blue bars).

MediFor MFC18 VGMD-IO CD@FAR=0.05





Figure: MediFor Video GAN CD@FAR=0.05

Takeaway

- MediFor MFC19/MFC18 results, cross-year comparisons
- OpenMFC results
- Media Forensics is still in the early stages of development
- **Join the OpenMFC program! <https://mfc.nist.gov>**

Questions?

OpenMFC team: mfc_poc@nist.gov

The background of the slide is a light gray with a complex, abstract pattern of thin lines and dots, resembling a network or circuit board. There are also larger, faint geometric shapes like rectangles and circles scattered across the background.

Thank You!