

Deep Learning Innovations at the British Machine Vision Conference (BMVC)

Author: BenchChem Technical Support Team. **Date:** December 2025

Compound of Interest

Compound Name: *BMVC*

Cat. No.: *B3029348*

[Get Quote](#)

This document provides detailed application notes and protocols for key deep learning techniques presented at the British Machine Vision Conference (**BMVC**), tailored for researchers, scientists, and drug development professionals. The content summarizes novel methodologies, presents quantitative data in structured tables for comparative analysis, and includes detailed experimental protocols. Signaling pathways, experimental workflows, and logical relationships are visualized using Graphviz diagrams to facilitate understanding.

MeTTA: Single-View to 3D Textured Mesh Reconstruction with Test-Time Adaptation

This work introduces a novel test-time adaptation (TTA) method, named MeTTA, for reconstructing 3D textured meshes from a single image. This technique is particularly effective for out-of-distribution (OoD) samples, where traditional learning-based models often fail. By leveraging a generative prior and jointly optimizing 3D geometry, appearance, and pose, MeTTA can adapt to unseen objects at test time. A key innovation is the use of learnable virtual cameras with self-calibration to resolve ambiguities in the alignment between the reference image and the 3D shape.^{[1][2]}

Experimental Protocols

The core of MeTTA's methodology lies in its test-time adaptation pipeline which refines an initial coarse 3D model.

- **Initial Reconstruction:** A pre-trained feed-forward model provides an initial prediction of the 3D mesh and viewpoint from a single input image.
- **Test-Time Adaptation:** The initial prediction is then refined through an optimization process that minimizes a loss function combining a multi-view diffusion prior, a segmentation loss, and a regularization term.
- **Joint Optimization:** The optimization is performed jointly on the 3D mesh vertices, physically-based rendering (PBR) texture properties, and virtual camera parameters.
- **Learnable Virtual Cameras:** To handle potential misalignments, learnable virtual cameras are introduced. These cameras are optimized to find the best possible alignment between the rendered 3D model and the input image.
- **Generative Prior:** A pre-trained multi-view generative model (Zero-1-to-3) is used as a prior to guide the reconstruction, ensuring plausible 3D shapes even from a single view.[\[3\]](#)

The implementation details are as follows:

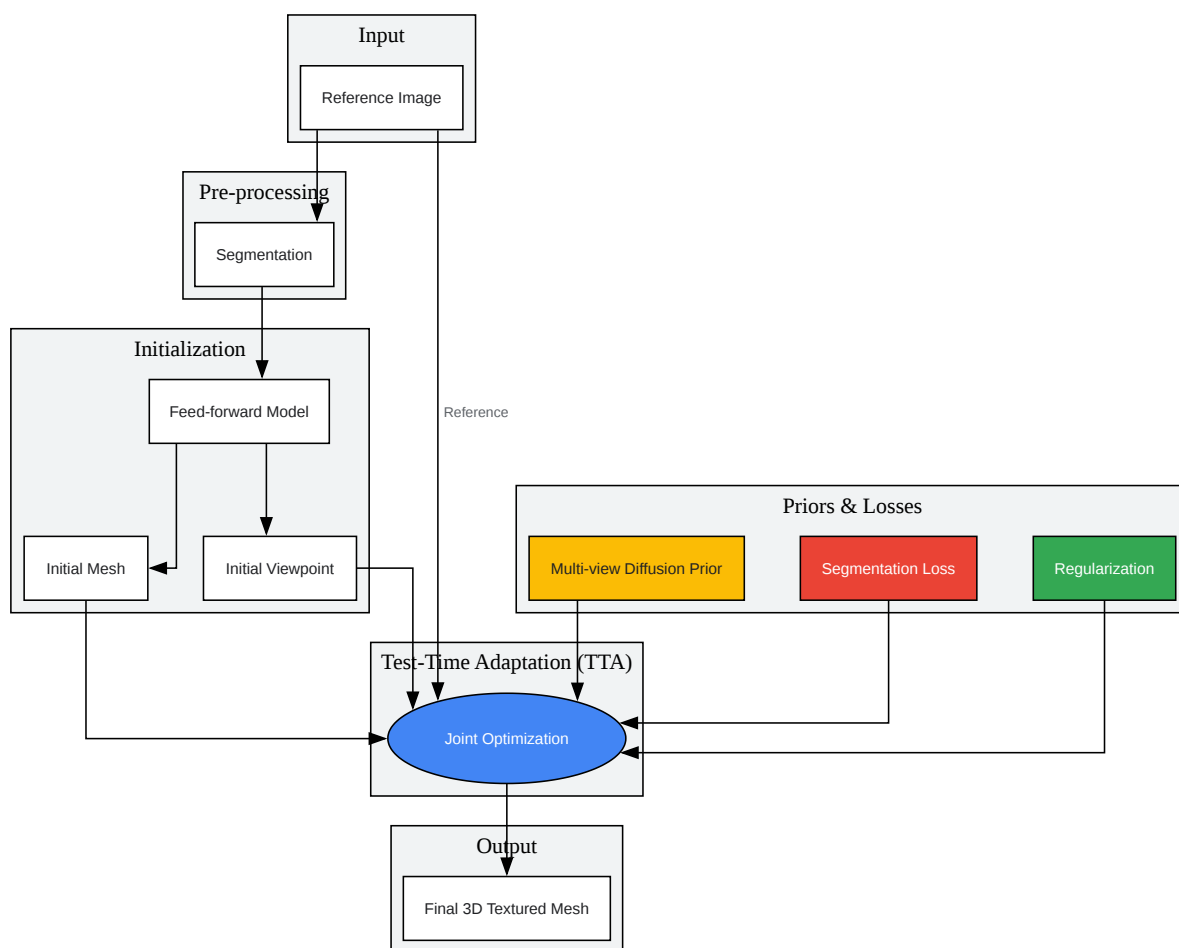
- **Environment:** The system requires an NVIDIA GPU with at least 48GB of VRAM for the default settings. The software stack includes Python 3.9, PyTorch, and other dependencies as specified in the official repository.[\[3\]](#)
- **Pre-processing:** Input images are first segmented to isolate the object of interest. The authors use the Grounded-Segment-Anything model for this purpose.[\[3\]](#)
- **Optimization:** The test-time adaptation for a single model takes approximately 30 minutes for 1500 iterations with a batch size of 8 on a single A6000 GPU.[\[3\]](#)

Data Presentation

The effectiveness of MeTTA was demonstrated through qualitative results on in-the-wild images where existing methods failed. The visual improvements in geometry and texture realism for out-of-distribution objects are the primary quantitative evidence of the method's success. The paper also includes ablation studies to validate the contribution of each component of the pipeline.[\[1\]](#)

Component	Contribution
Initial Mesh Prediction	Provides a starting point for the optimization, without which the model would have to start from a simple ellipsoid, leading to slower convergence and potentially poorer results.
Initial Viewpoint Prediction	Crucial for establishing an initial alignment. Without it, the optimization starts from a canonical viewpoint, which may be far from the correct one.
Learnable Virtual Cameras	Significantly improves the alignment between the rendered model and the reference image, correcting for errors in the initial viewpoint prediction.
Multi-view Diffusion Prior	Enforces a strong 3D shape prior, preventing unrealistic geometries and guiding the reconstruction towards plausible shapes.

Mandatory Visualization



[Click to download full resolution via product page](#)

Caption: The MeTTA pipeline for single-view 3D reconstruction with test-time adaptation.

FedFS: Federated Learning for Face Recognition via Intra-subject Self-supervised Learning

This paper proposes a novel federated learning framework, FedFS, designed for personalized face recognition.^[4] It addresses two key challenges in existing federated learning approaches for face recognition: the insufficient use of self-supervised learning and the requirement for clients to have data from multiple subjects.^[4] FedFS enables the training of personalized face recognition models on devices with data from only a single subject, enhancing data privacy.^[5]

Experimental Protocols

The FedFS framework consists of two main components that work in conjunction with a pre-trained feature extractor.

- **Adaptive Soft Label Construction:** This component reformats labels within intra-instances using dot product operations between features from the local model, the global model, and an off-the-shelf pre-trained model. This allows the model to learn discriminative features for a single subject.
- **Intra-subject Self-supervised Learning:** Cosine similarity operations are employed to enforce robust intra-subject representations. This helps in reducing the intra-class variation for the features of a single individual.
- **Regularization Loss:** A regularization term is introduced to prevent the personalized model from overfitting to the local data and to ensure the stability of the optimized model.^[4]

The experimental setup is as follows:

- **Datasets:** The effectiveness of FedFS was evaluated on the DigiFace-1M and VGGFace datasets.^[4]
- **Pre-trained Model:** The PocketNet model was used as the off-the-shelf pre-trained feature extractor.^[6]
- **Federated Learning Setup:** The participation rate of clients in each round of federated learning was set to 0.7.^[6]

- Optimization: The local models were trained using the proposed loss function, and the global model was updated by aggregating the parameters of the participating local models.

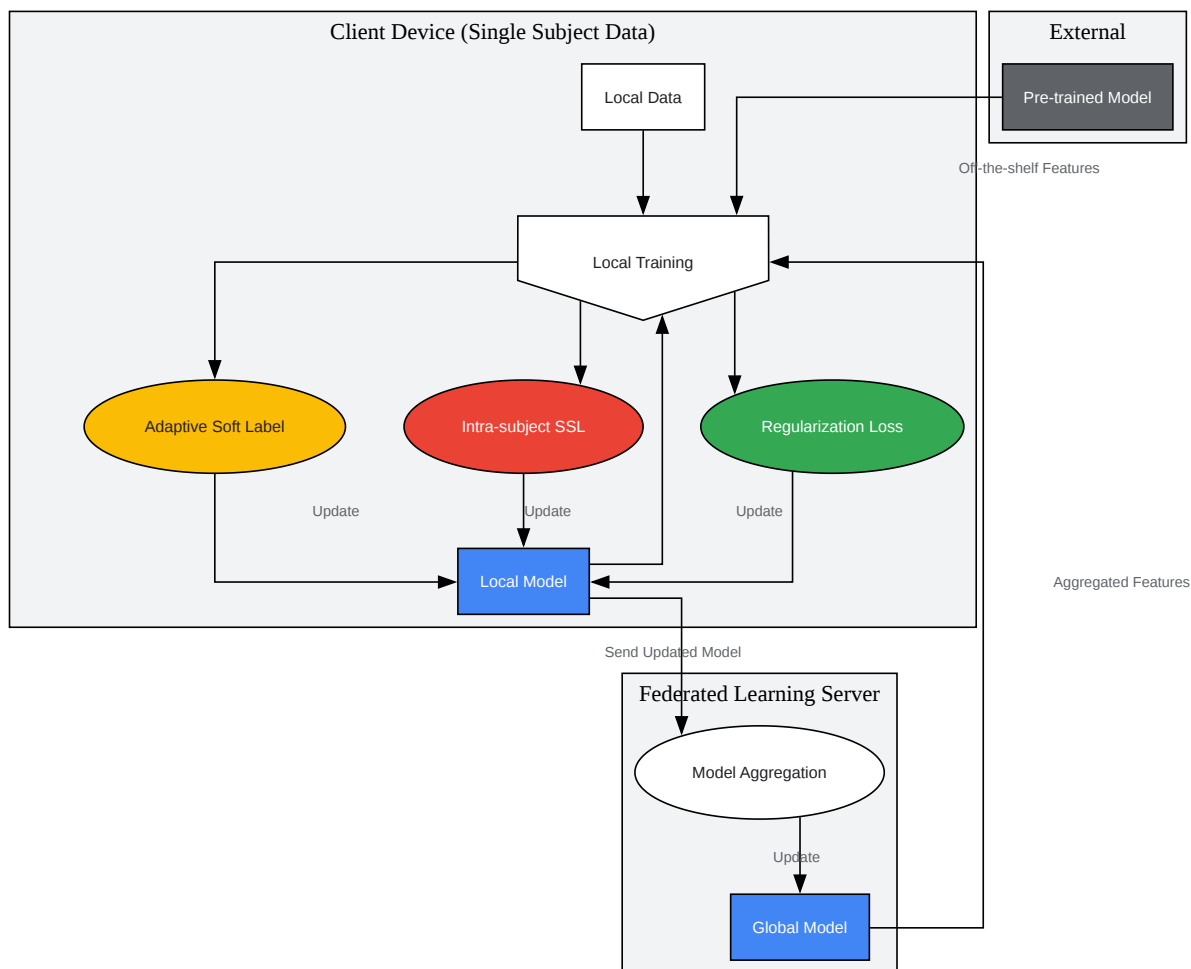
Data Presentation

The performance of FedFS was compared against previous methods, demonstrating superior performance. The key results are summarized below.

Method	Dataset	Metric	Performance
Previous SOTA	DigiFace-1M	Accuracy	Lower
FedFS	DigiFace-1M	Accuracy	Higher
Previous SOTA	VGGFace	Accuracy	Lower
FedFS	VGGFace	Accuracy	Higher

Furthermore, analysis showed that the intra-subject self-supervised learning component effectively reduces intra-class variance, as indicated by a smaller intersection of positive and negative similarity areas in the feature space.[\[6\]](#)

Mandatory Visualization



[Click to download full resolution via product page](#)

Caption: The FedFS framework for personalized federated face recognition.

Efficiency-preserving Scene-adaptive Object Detection

This work tackles the problem of adapting object detection models to new scenes without the need for manual annotation. It proposes a self-supervised scene adaptation framework that is also efficiency-preserving. This is an extension of their previous work on object detection with self-supervised scene adaptation presented at CVPR 2023.[\[7\]](#)

Experimental Protocols

The proposed framework enables a pre-trained object detector to adapt to a new target scene using only unlabeled video frames from that scene.

- **Self-Supervised Learning:** The core of the method is a self-supervised learning approach where pseudo-labels are generated for the unlabeled target scene data.
- **Fusion Network:** A key component is a fusion network that takes object masks as an additional input modality to the standard RGB input. This helps the model to better distinguish objects from the background.
- **Dynamic Background Generation:** To improve the robustness of the model, dynamic background images are generated from the video frames. This is achieved by using image inpainting techniques.
- **Pseudo-Label Generation and Refinement:** The initial pseudo-labels are generated by a pre-trained model and then refined using a graph-based method.
- **Adaptation Training:** The object detection model is then fine-tuned on the target scene data using the refined pseudo-labels.

The implementation details are as follows:

- **Environment:** The system requires an NVIDIA GPU with at least 20GB of VRAM. It is built upon Detectron2 v0.6.[\[7\]](#)
- **Dataset:** The paper introduces the Scenes100 dataset for evaluating scene-adaptive object detection.[\[7\]](#)

- Training: The adaptation training is performed using a provided script, and the resulting checkpoints are saved for evaluation.

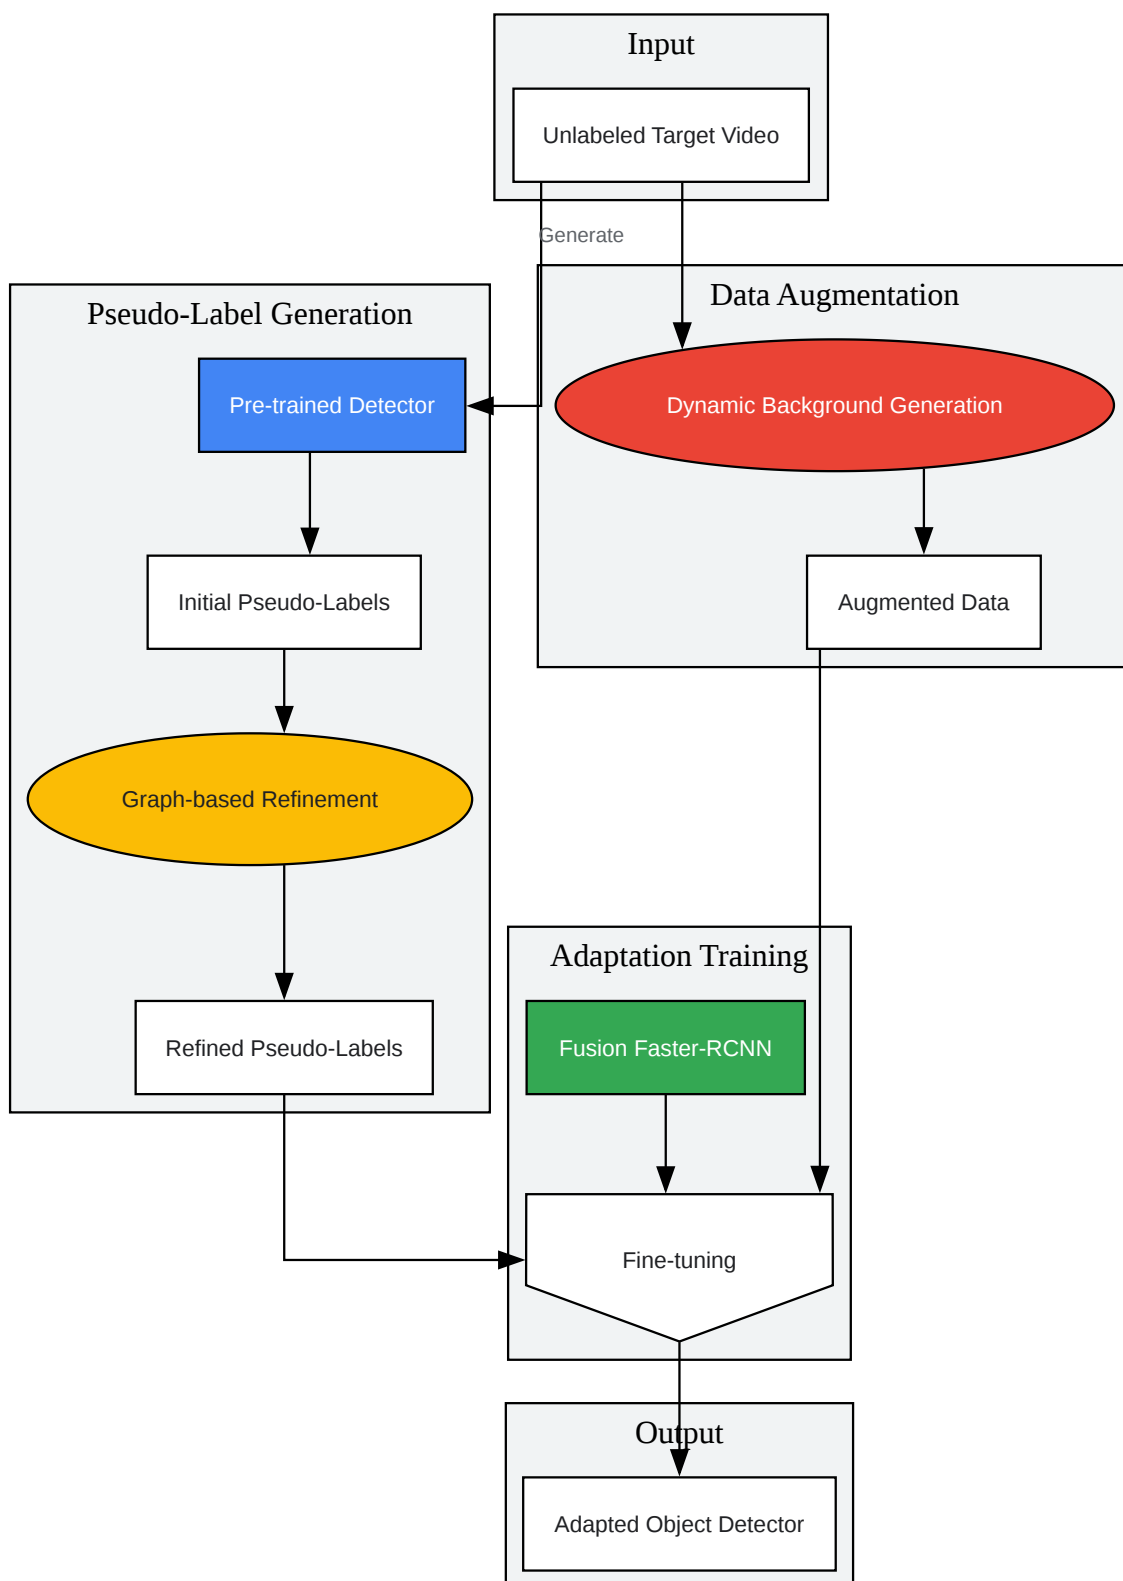
Data Presentation

The quantitative results from the associated CVPR 2023 paper, which this work extends, demonstrate the effectiveness of the self-supervised adaptation. The performance is measured in terms of Average Precision (AP) on the Scenes100 dataset.

Method	Adaptation	AP on Scenes100
Baseline (no adaptation)	No	Lower
Self-Supervised Adaptation	Yes	Higher

The **BMVC** 2024 paper focuses on making this adaptation process more efficient.

Mandatory Visualization



[Click to download full resolution via product page](#)

Caption: Workflow for efficiency-preserving scene-adaptive object detection.

Need Custom Synthesis?

BenchChem offers custom synthesis for rare earth carbides and specific isotopic labeling.

Email: info@benchchem.com or [Request Quote Online](#).

References

- 1. bmva-archive.org.uk [bmva-archive.org.uk]
- 2. [2408.11465] MeTTA: Single-View to 3D Textured Mesh Reconstruction with Test-Time Adaptation [arxiv.org]
- 3. GitHub - kaist-ami/MeTTA: [BMVC'24] Official repository for "MeTTA: Single-View to 3D Textured Mesh Reconstruction with Test-Time Adaptation" [github.com]
- 4. [2407.16289] Federated Learning for Face Recognition via Intra-subject Self-supervised Learning [arxiv.org]
- 5. bmva-archive.org.uk [bmva-archive.org.uk]
- 6. bmva-archive.org.uk [bmva-archive.org.uk]
- 7. GitHub - cvlab-stonybrook/scenes100: Official repo of Object Detection with Self-Supervised Scene Adaptation [CVPR 2023] and Efficiency-preserving Scene-adaptive Object Detection [BMVC 2024] [github.com]
- To cite this document: BenchChem. [Deep Learning Innovations at the British Machine Vision Conference (BMVC)]. BenchChem, [2025]. [Online PDF]. Available at: [<https://www.benchchem.com/product/b3029348#deep-learning-techniques-presented-at-bmvc>]

Disclaimer & Data Validity:

The information provided in this document is for Research Use Only (RUO) and is strictly not intended for diagnostic or therapeutic procedures. While BenchChem strives to provide accurate protocols, we make no warranties, express or implied, regarding the fitness of this product for every specific experimental setup.

Technical Support: The protocols provided are for reference purposes. Unsure if this reagent suits your experiment? [[Contact our Ph.D. Support Team for a compatibility check](#)]

Need Industrial/Bulk Grade? [Request Custom Synthesis Quote](#)

BenchChem

Our mission is to be the trusted global source of essential and advanced chemicals, empowering scientists and researchers to drive progress in science and industry.

Contact

Address: 3281 E Guasti Rd
Ontario, CA 91761, United States
Phone: (601) 213-4426
Email: info@benchchem.com