EUROPEAN CONFERENCE ON COMPUTER VISION

MILANO 2024

# CriSp: Leveraging Tread Depth Maps for Enhanced Crime-Scene Shoeprint Matching

Samia Shafique[1], Shu Kong[2]*, and Charless Fowlkes[1]*

[1]University of California, Irvine, [2]Texas A&M University

# Crime-scene Shoeprints

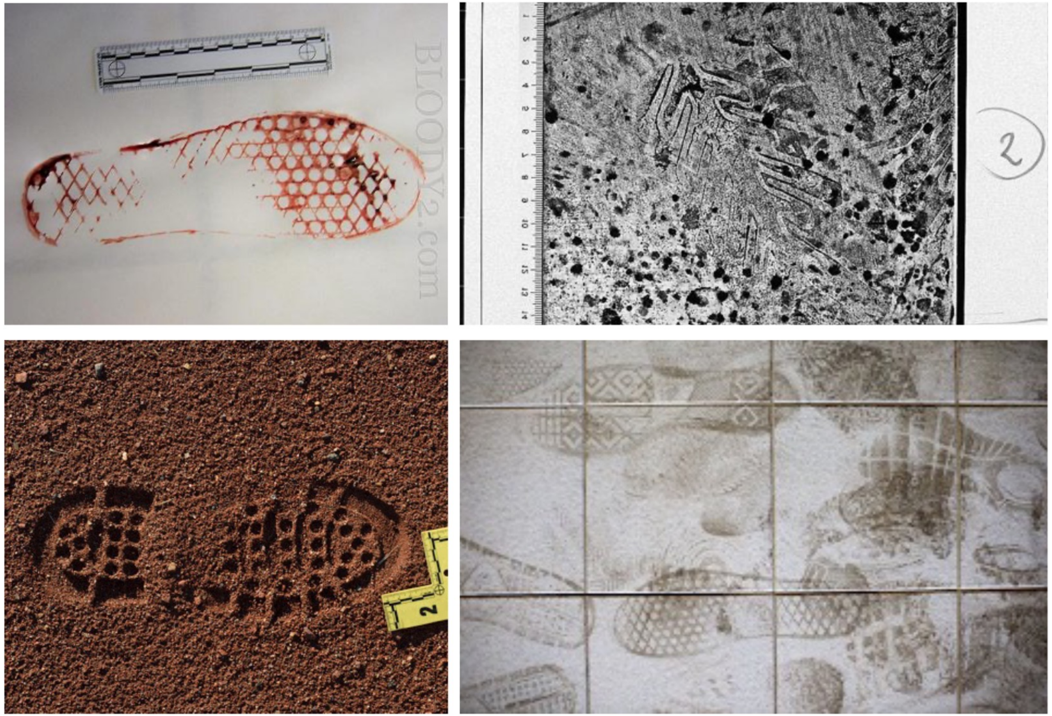


- Examining evidence from a crime scene helps investigators identify suspects.
- Challenge in analyzing
  - Noisy and degraded
  - Occluded
  - Partial
  - Occur across various mediums

# Matching Print to Shoe

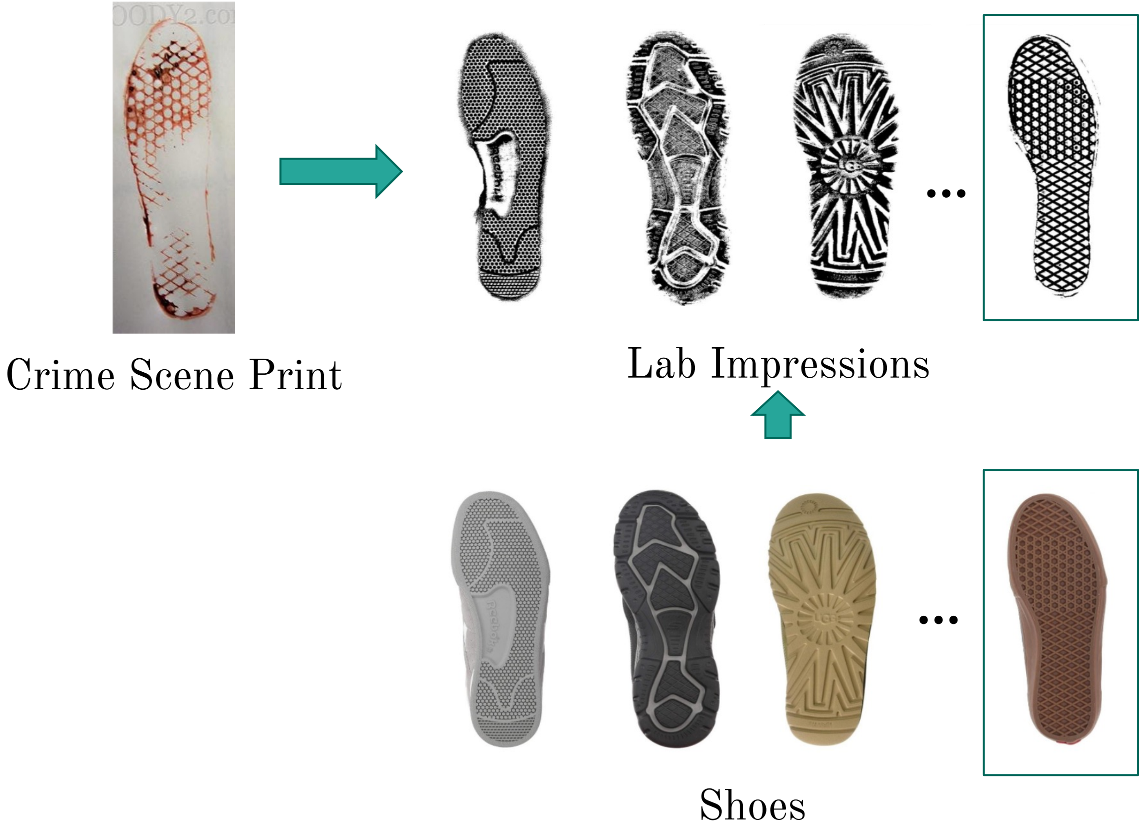


Previous work

- Compare crime-scene shoeprints to prints in small, manually curated reference database.

# Matching Print to Shoe

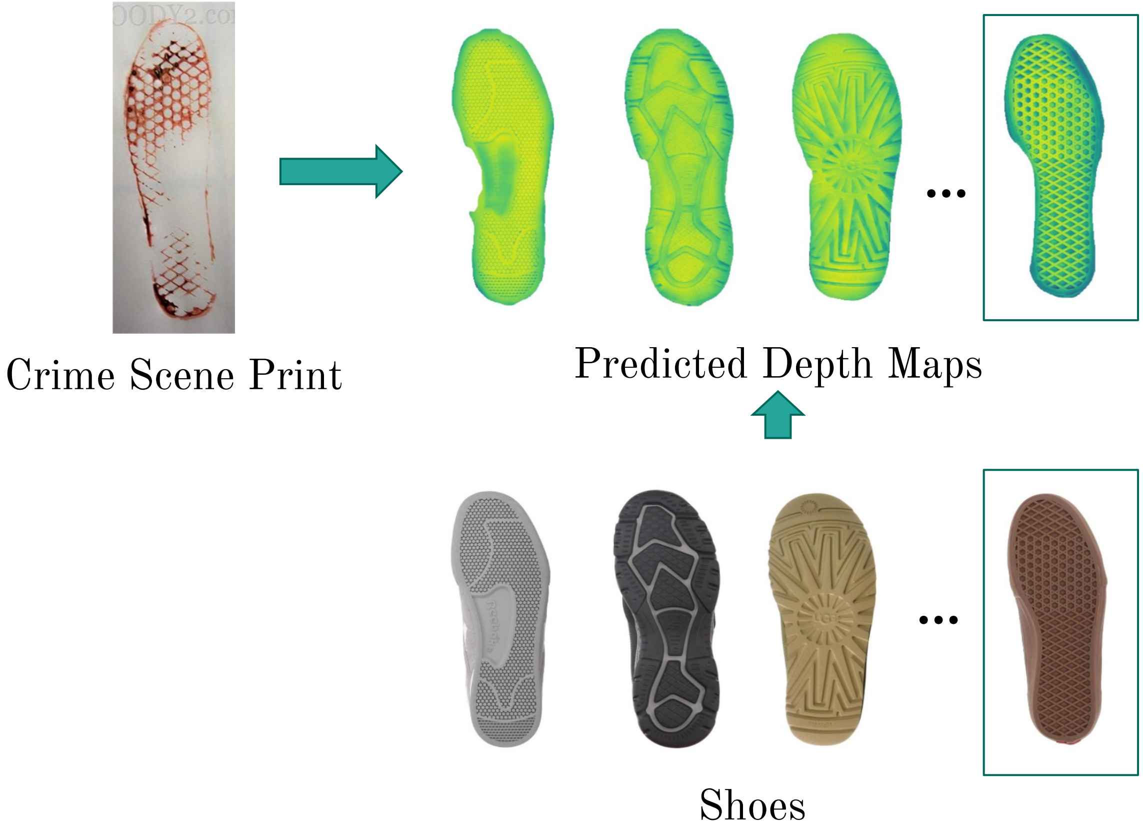


Previous work

- Compare crime-scene shoeprints to prints in small, manually curated reference database.

CriSp

- Compare crime-scene shoeprints to depth maps from large-scale, automatically generated reference database of tread images from online retailers.

# Training and Validation Datasets

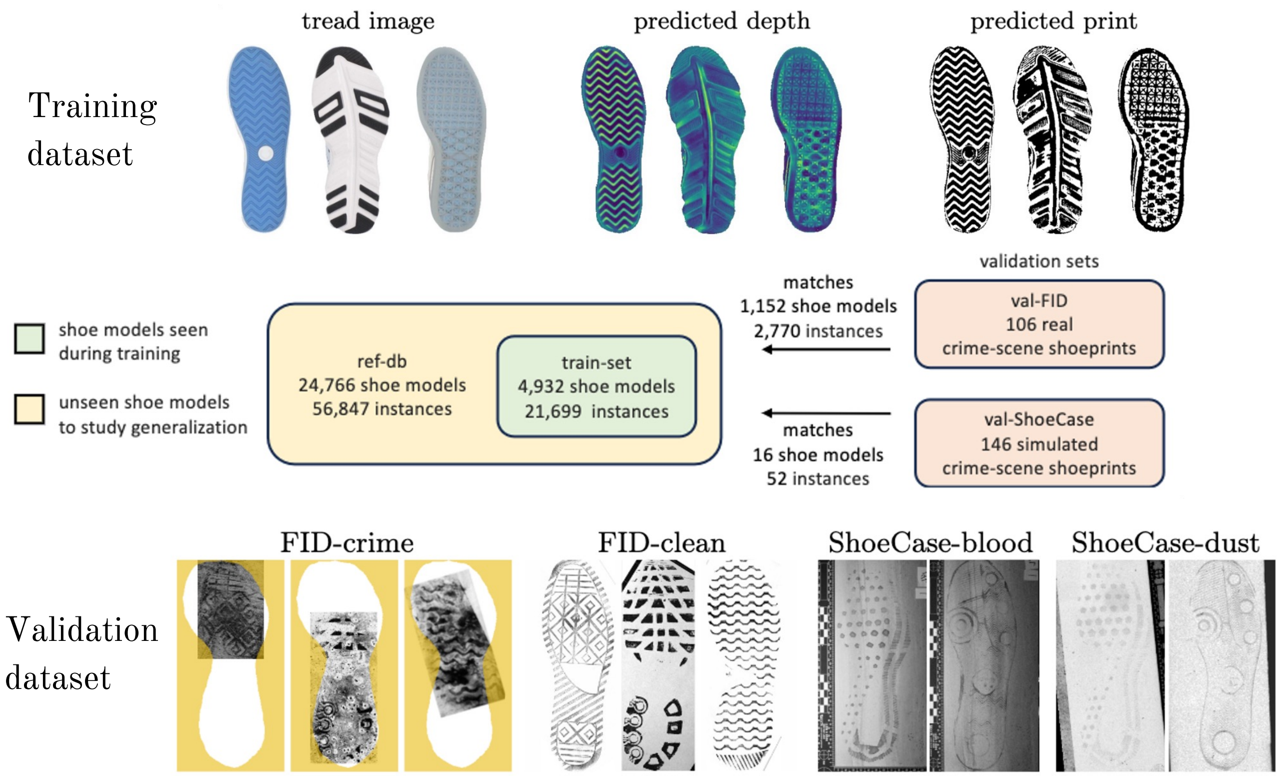


Training data: Tread images with *predicted* depth maps and prints.

Reference Database: Superset of training data.

Validation sets:

- val-FID: real crime-scene shoeprints from FID300 [1].
- val-ShoeCase: simulated crime-scene shoeprints from ShoeCase [2].


[1] Kortylewski et al. Unsupervised Footwear Impression Analysis and Retrieval from Crime Scene Data. In ACCVW 2014.
[2] Tibben et al. ShoeCase: A data set of mock crime scene footwear impressions. In Data in Brief 2023.

# Methodology

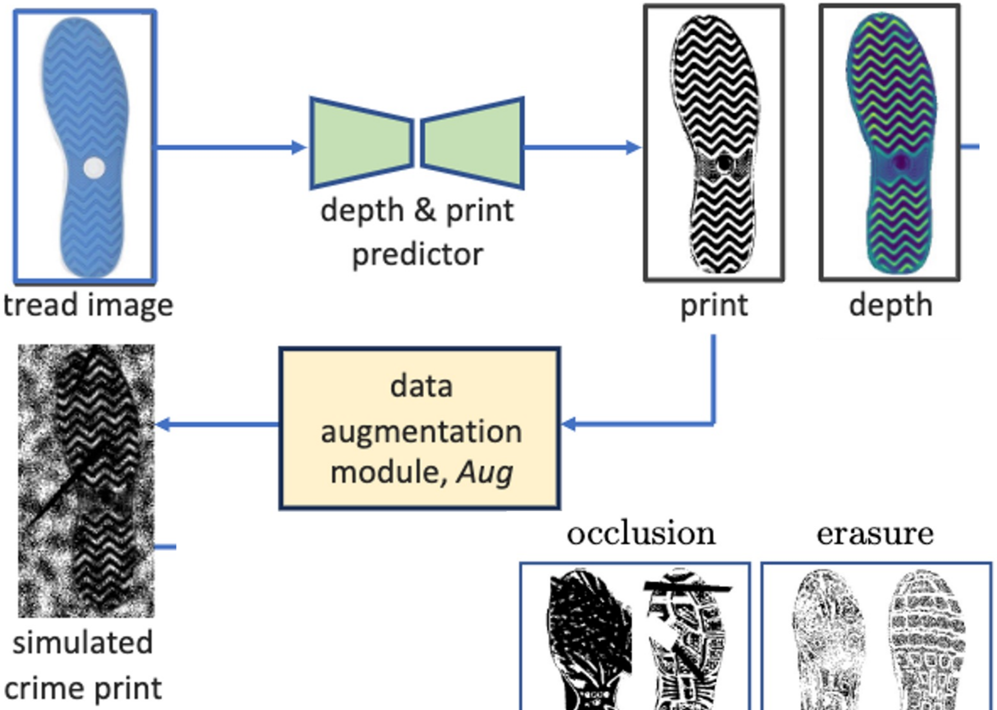


- Data augmentation simulates crime scene shoeprints.
  - Occlusion - overlapping prints and quads
  - Erasure - grainy nature of crime-scene prints
  - Noise - background clutter

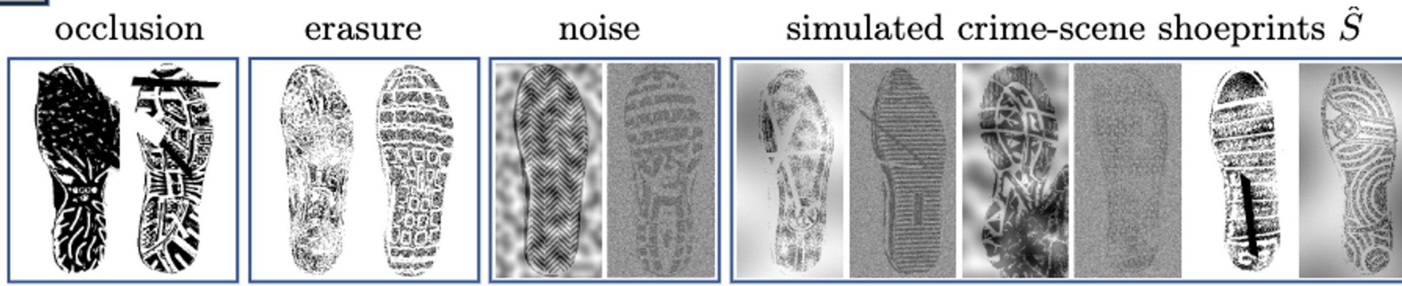

# Methodology

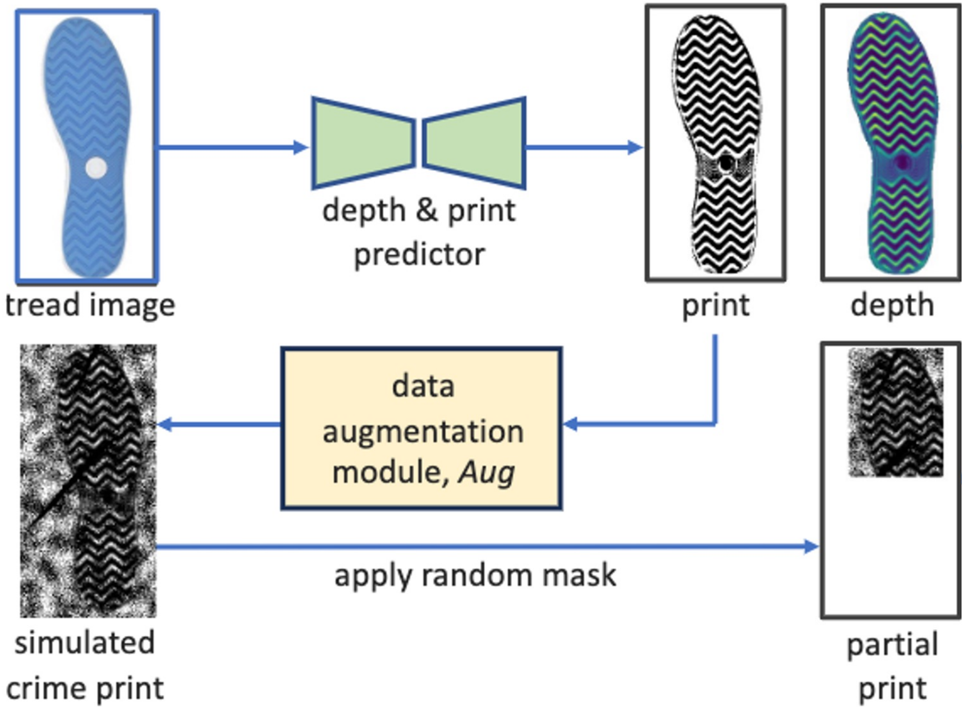


- Generate partial prints by applying a random mask on the simulated crime scene prints.

# Methodology

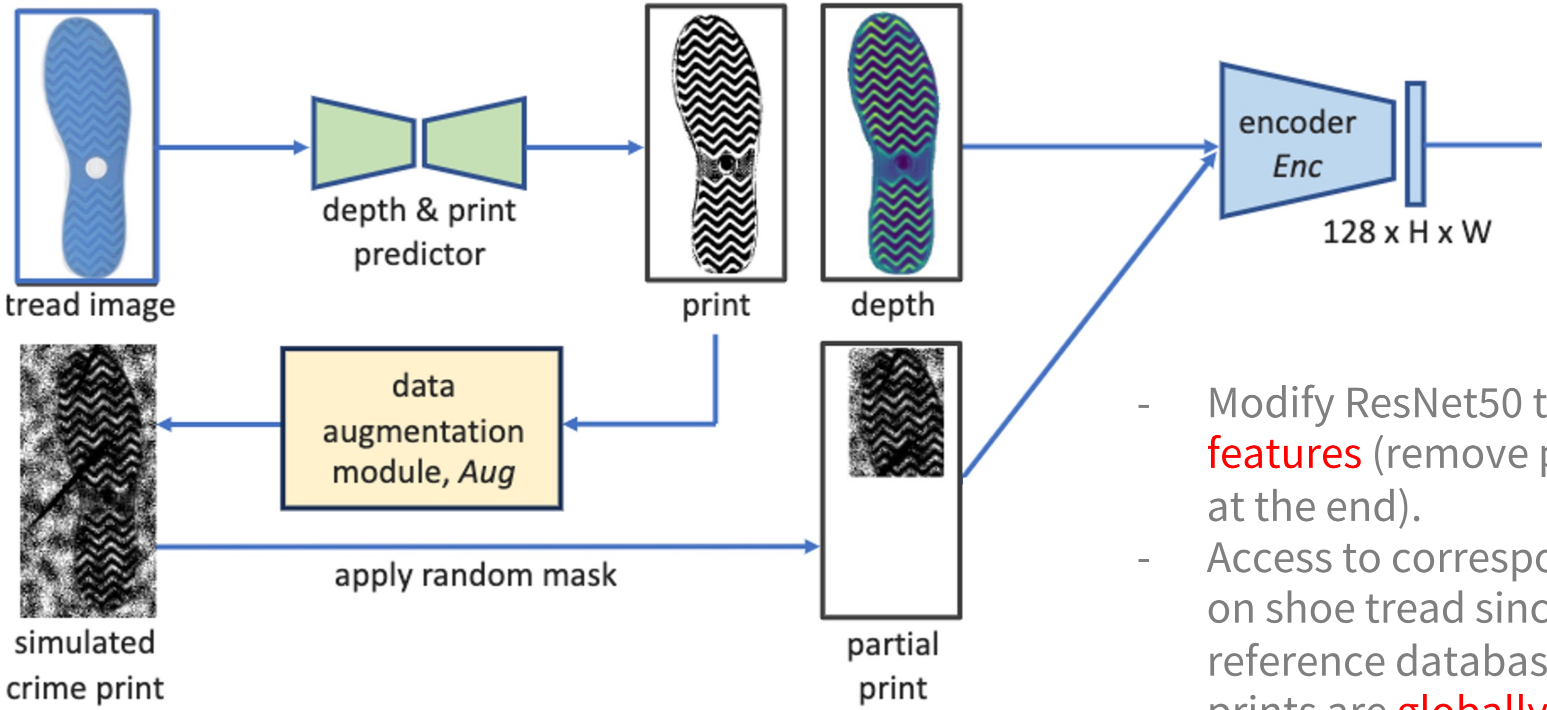


- Modify ResNet50 to output spatial features (remove pooling operation at the end).
- Access to corresponding locations on shoe tread since training data, reference database and validation prints are globally aligned.

# Methodology

- Use a compressed mask of size HxW to mask out irrelevant portions of the spatial features.

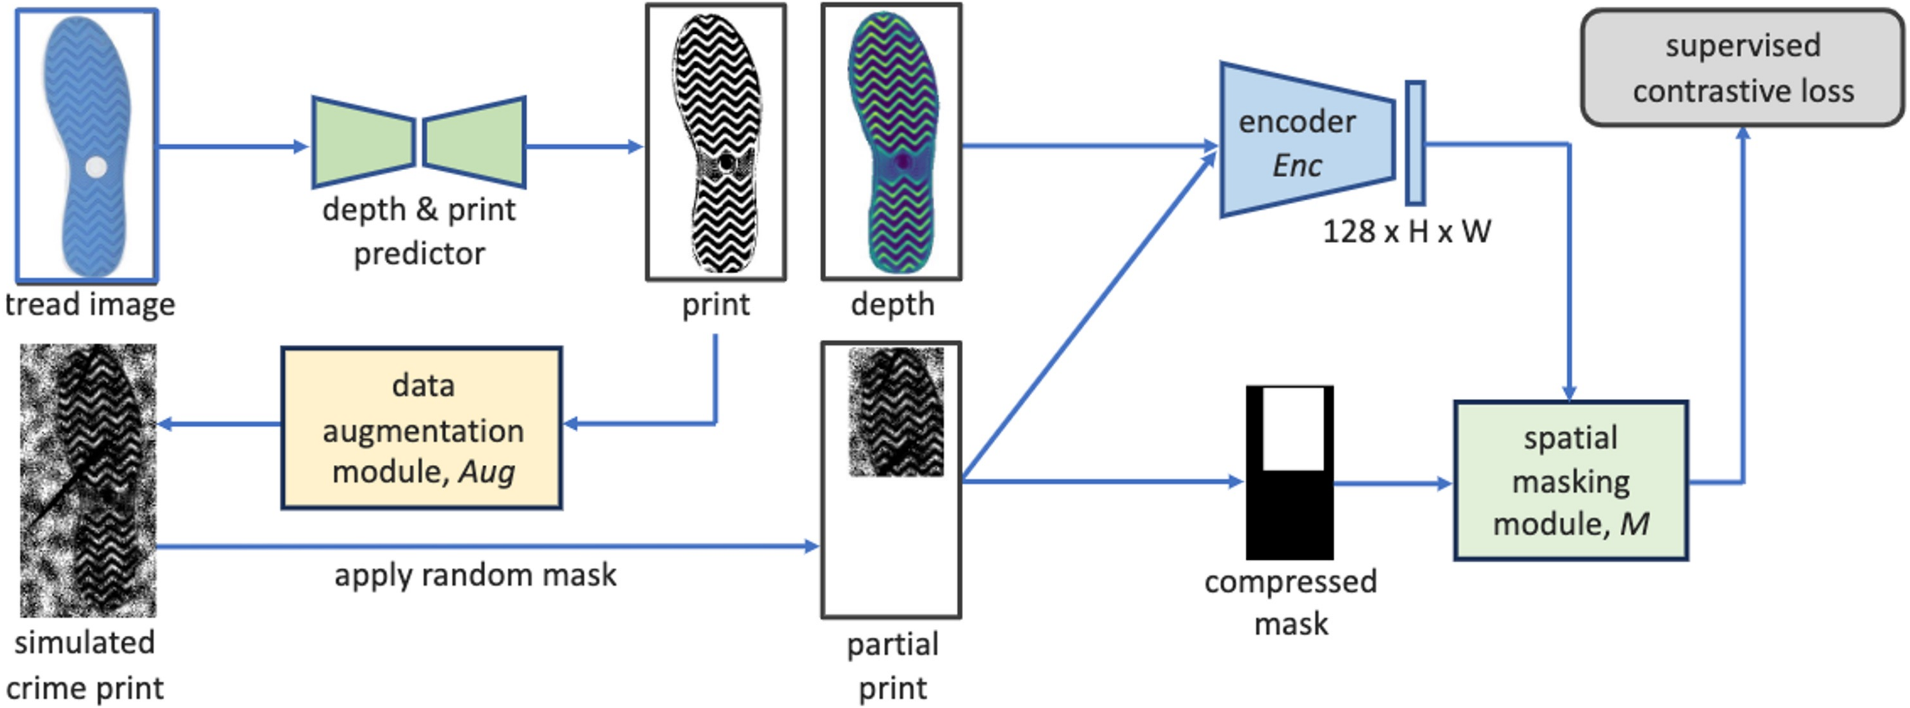

# Qualitative Results

We visualize the top 10 retrievals of CriSp on val-FID (row1-5) and val-ShoeCase (row 6).

CriSp retrieves positives matches early even with very limited visibility or severe degradation.

# Quantitative Comparison to SOTA

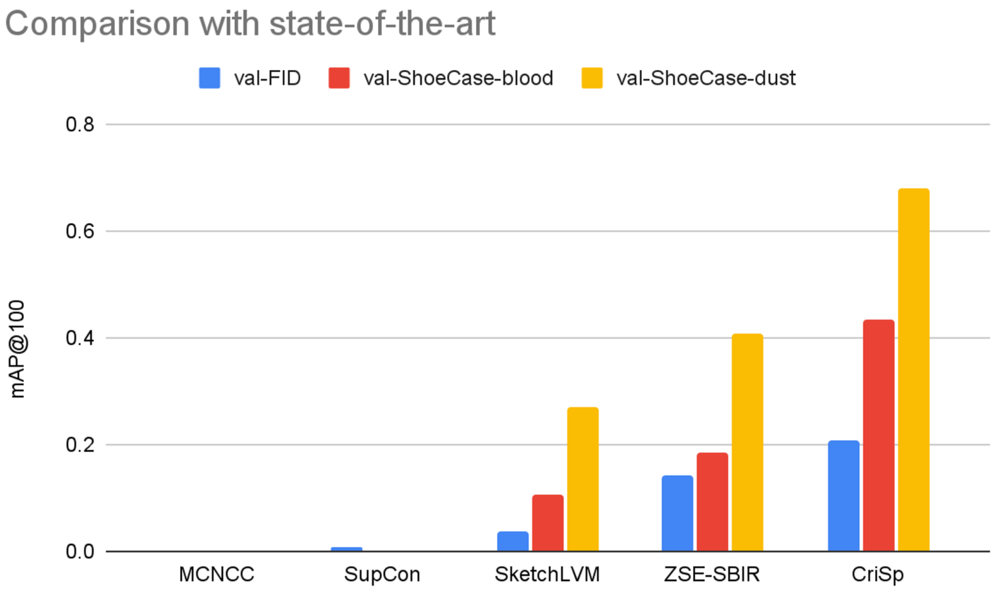


We compare to SOTA on

- automated shoeprint matching (MCNCC) and
- image retrieval tailored to this task (SupCon, SketchLVM, ZSE-SBIR)

MCNCC: Kong et al. Cross-domain image matching with deep feature maps. IJCV 2019
SupCon: Khosla et al. Supervised contrastive learning. In NeurIPS 2020
SketchLVM: Sain et al. Clip for all things zero-shot sketch-based image retrieval, fine-grained or not. CVPR 2023.
ZSE-SBIR: Lin et al. Zero-shot everything sketch-based image retrieval, and in explainable style. CVPR 2023.

# Thank You!