F²DC

Federated learning that generalizes across domains

Feature Decoupling and Calibration for domain-skewed federated learning.

Key message: turn domain bias from noise into calibrated evidence.

Huan Wang, Jun Shen, Jun Yan, Guansong Pang

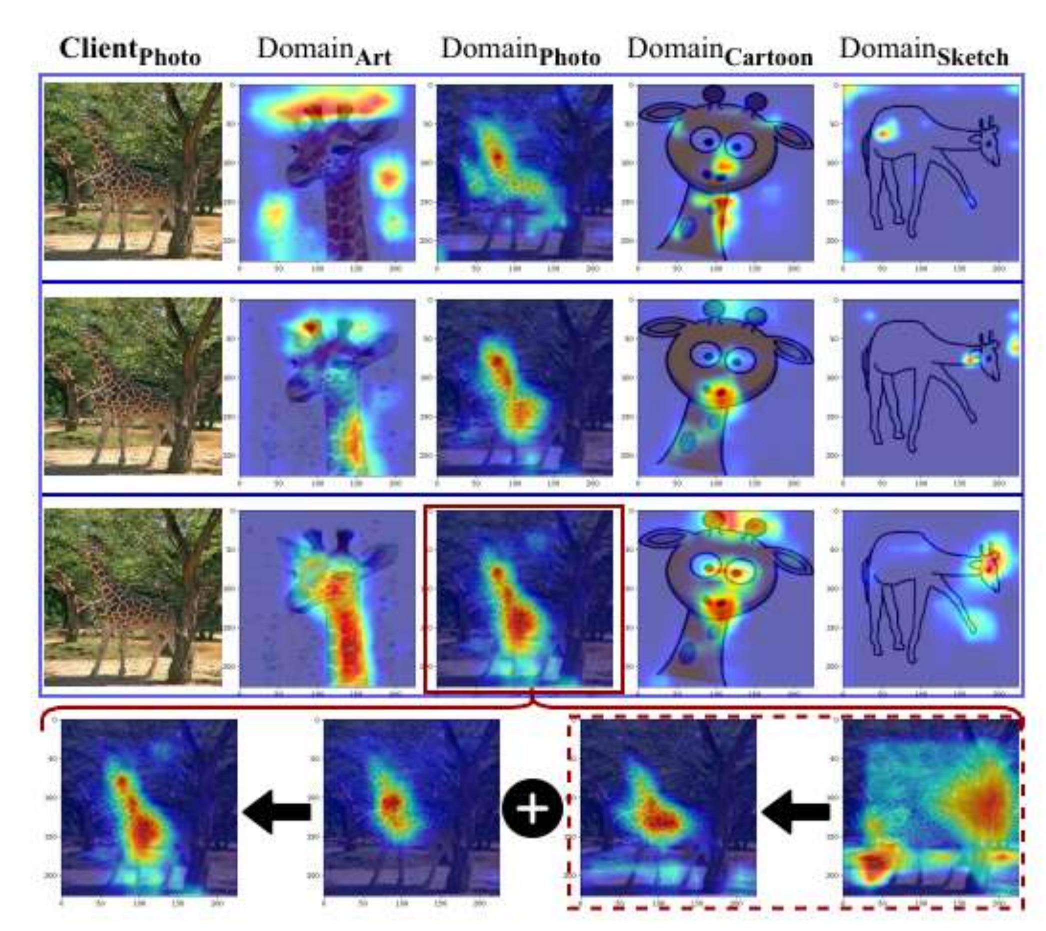

visual domains → biased local features

The challenge: clients see different “styles”

Federated learning usually assumes local updates point toward a compatible global model. Domain skew breaks this assumption.

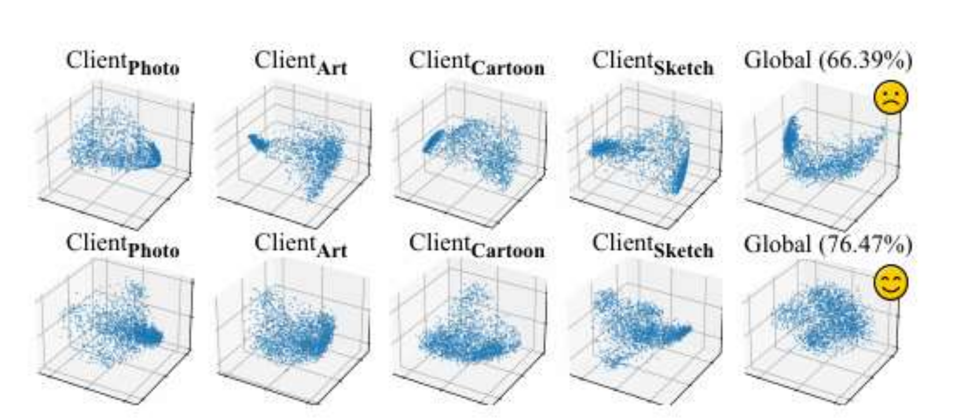

Representations collapse into narrow, client-specific subspaces.

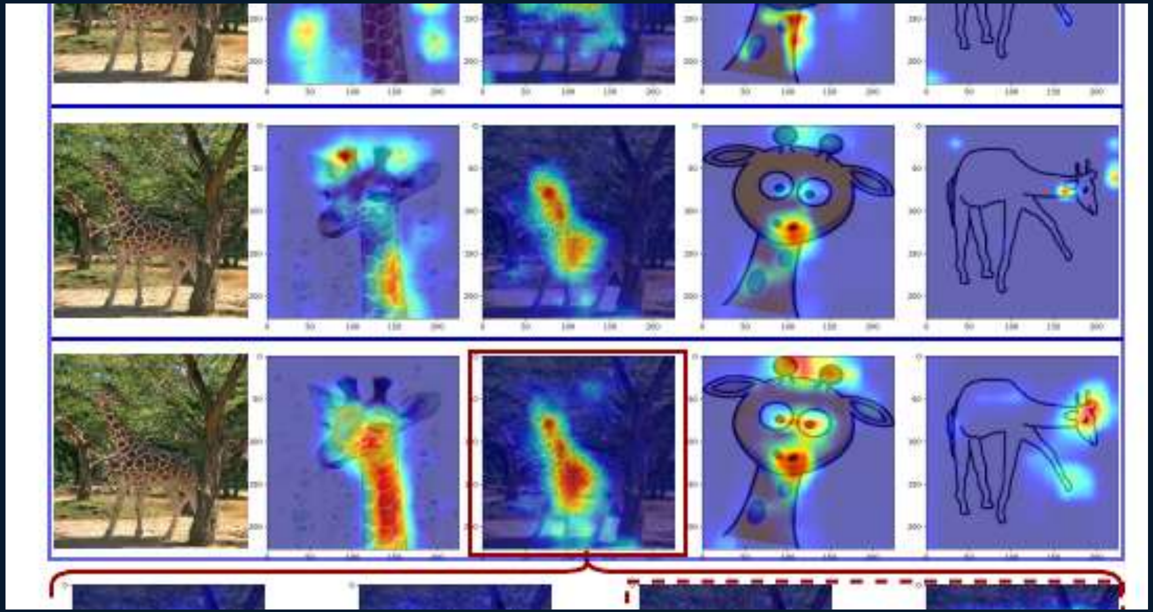

Attention maps become inconsistent across domains.

Plain-language view

A local model may learn “this looks like a sketch giraffe” instead of “this is a giraffe”. Aggregating such biased models weakens the global model.

Core insight: bias may still contain class evidence

Prior intuition

Domain-related features → Erase

Risk: useful class clues are removed together with nuisance style.

F^2DC intuition

- **1 · Decouple** — separate robust evidence from domain-related evidence
- **2 · Correct** — recover useful semantics hidden inside biased features
- **3 · Aggregate** — weight clients using data size and domain discrepancy

One-sentence summary

F^2DC does not just “de-bias”; it converts biased local features into calibrated, transferable evidence.

Method: three moves, one global objective

The model keeps raw data local. Extra modules are local too; only the standard model weights are sent to the server.

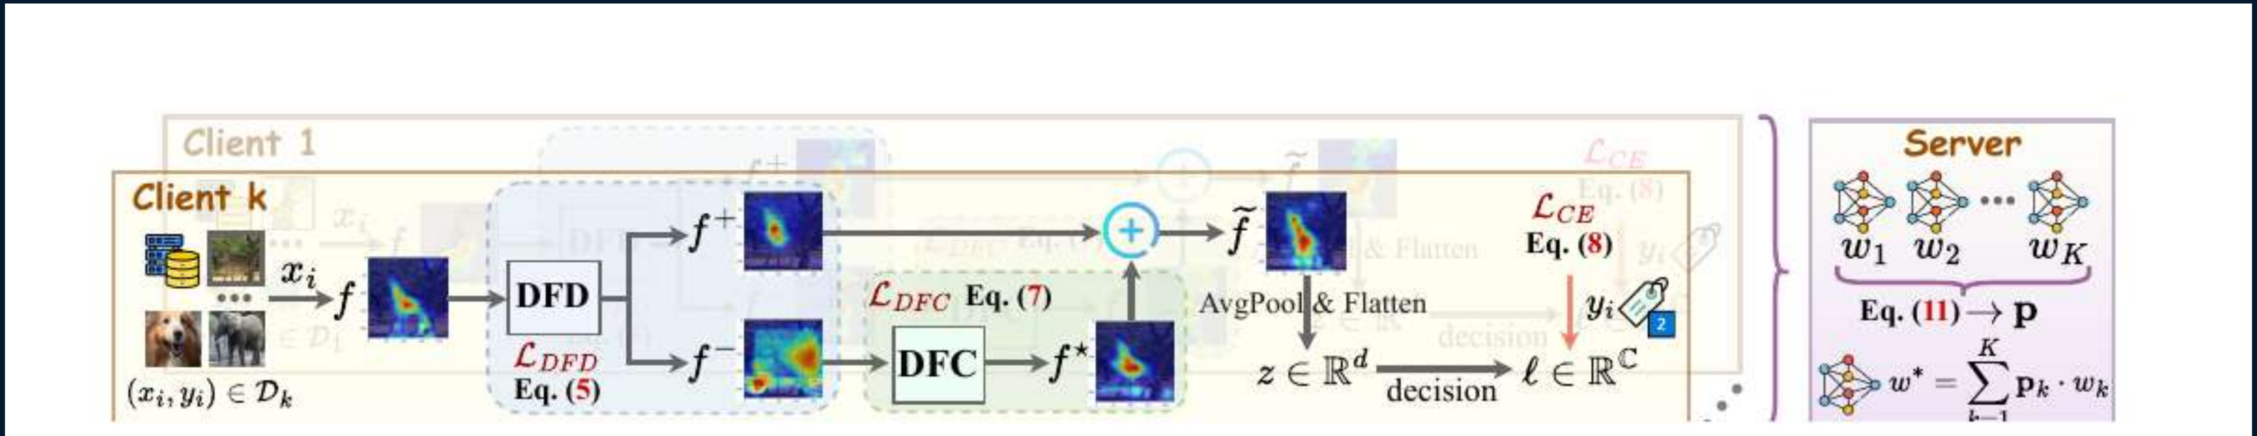

DFD
Decouple

Find which feature units are domain-robust and which are domain-related.

DFC
Correct

Calibrate domain-related features to recover class-relevant clues.

DaA
Aggregate

Use data size and domain discrepancy when forming the global model.

What changes after calibration?

The learned representation becomes less tied to one domain and more tied to class semantics.

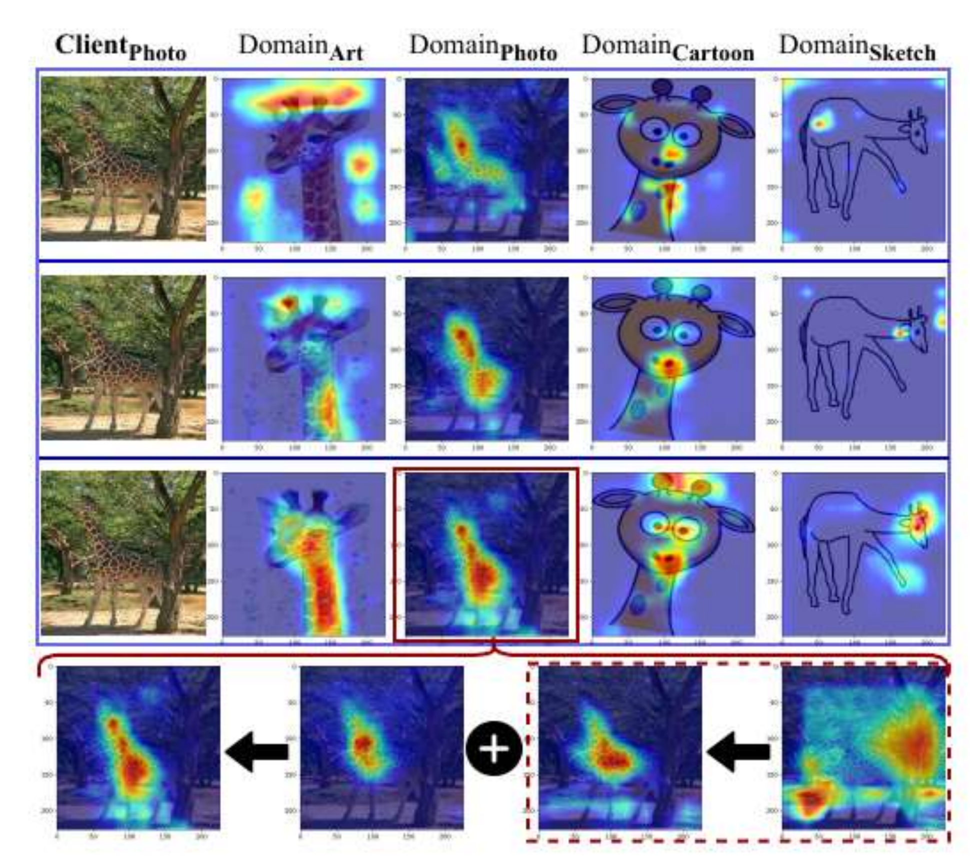

Attention evidence

Compared with vanilla FL and elimination-based baselines, F^2DC attends to more complete object evidence across domains.

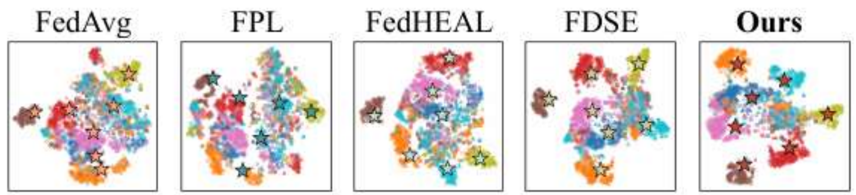

Figure 4. **T-SNE visualization** on PACS. Each color means one class, each shape means one domain, stars are semantic centers.

Representation space

Features of the same class form tighter clusters, while different classes separate more clearly.

Interpretation: the model learns "what the object is", not just "what the local domain looks like".

Results: higher accuracy, smaller domain gaps

Average top-1 accuracy improves across all three multi-domain benchmarks.

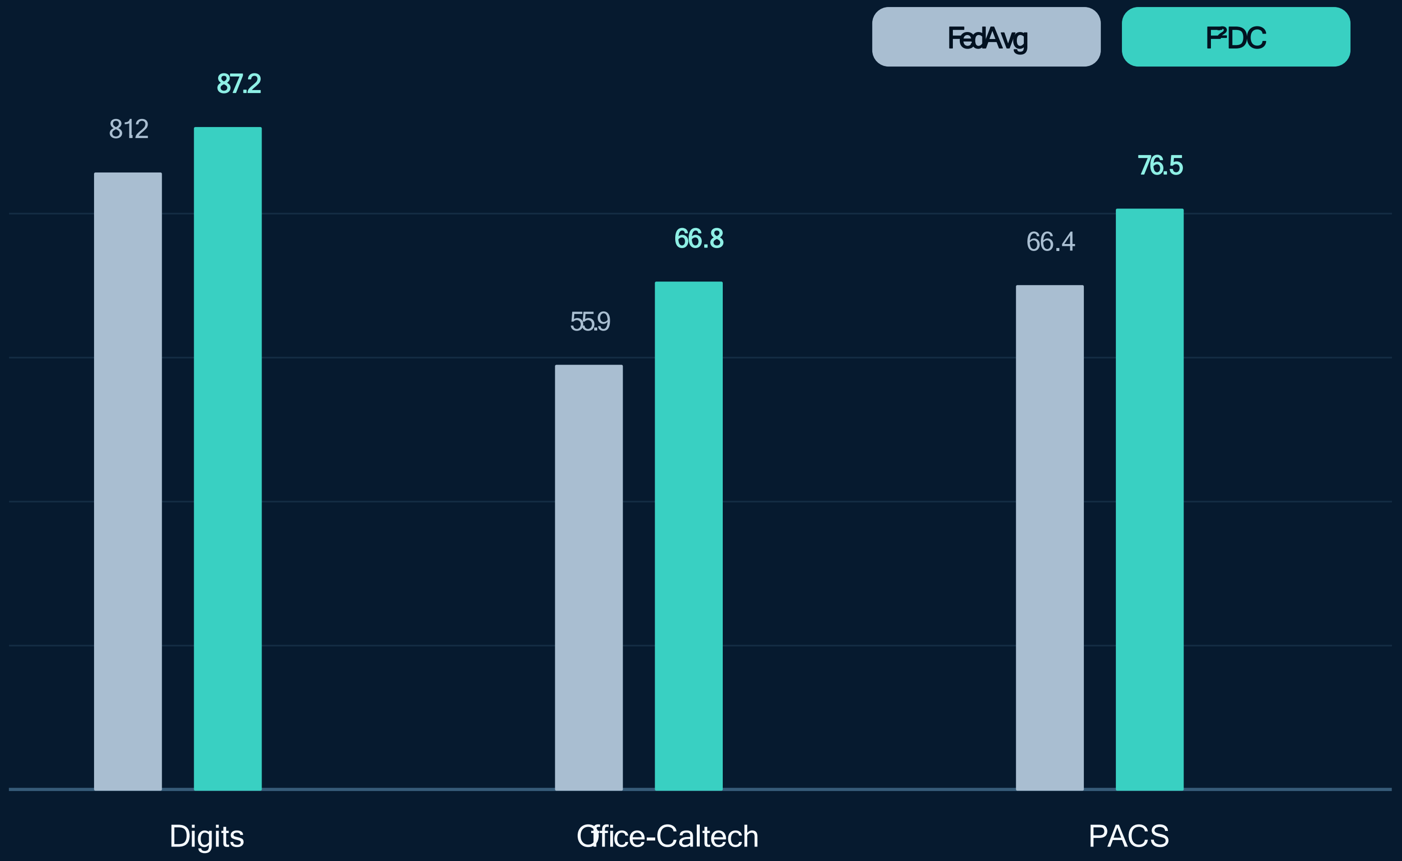

PACS
66.39 → 76.47

Office-Caltech
55.86 → 66.82

No extra communication
193.72MB per round on PACS

STD also decreases on all datasets, indicating fairer cross-domain performance.

Takeaway

F^2DC makes domain-skewed federated learning more generalizable by repairing, not discarding, domain-related features.

Problem	Clients learn biased representations because each client sees one domain.
Insight	Domain-related features include nuisance style and useful class evidence.
Method	Decouple, correct, and aggregate in a domain-aware manner.
Outcome	Higher accuracy, smaller domain gaps, and no additional communication cost.

Final message: F^2DC turns biased local evidence into transferable global evidence.