

Design Patent Similarity with ArcFace Metric Learning


Hui-Ching Chuang · Yueh-Heng Wu

Department of Statistics, National Taipei University · hcchuang@gm.ntpu.edu.tw

Patent identity—not product category—is the right supervision for finding visually similar designs.


15-MINUTE RESEARCH TALK

Three retrieval examples show success and hard negatives




Example 1 · exact recovery

Query D0807117: all five retrieved views share the query patent ID.



Example 2 · exact recovery

Query D0806975: all five retrieved views share the query patent ID.




Example 3 · hard negative

Query D0808886: visually close panels rank highly despite different patent IDs.

Extreme cosine cases expose the identity boundary

(a) Same-patent success



(b) High-confidence hard negative






Figure 3: Extreme baseline cosine-similarity examples on the Standard Gallery. Each panel shows a query, its highest-similarity gallery item, and its lowest-similarity item. Panel (a) is a same-patent success; panel (b) is a high-confidence hard negative despite very high visual similarity.

A national-scale corpus makes the task realistic



Train / validation / retrieval split




The extended gallery is 8.7× larger than the standard gallery—closer to production retrieval.

PhotoLike filtering removes photographs before training

Representative photographic-style patents



Pixel-intensity entropy separates the domains



Line drawings are sparse and low-entropy; photographs contain richer shading and texture.

Threshold

$\tau = 0.724$
99% target recall

Flagged

28,717
candidates

Excluded

25,448
verified patents

Patent IDs become the metric-learning signal



Training


Scale 30 · angular margin 0.50 · 30 epochs · effective batch ≈ 512

Retrieval

Cosine search, then optional query expansion or graph diffusion

Resolution and backbone capacity both matter


BEST: 0.8193 mAP



DESIGN PATENT SIMILARITY WITH ARCFACE METRIC LEARNING · HUI-CHING CHUANG

The gain survives a much larger search gallery


Standard Gallery



Acc@1 = 0.9312

Acc@20 = 0.9906

Extended Gallery




Acc@1 = 0.8088


Acc@20 = 0.9711

The embedding learns identity—not just taxonomy

Across the design hierarchy




Within one subclass (D14/299)



Same-patent views cluster strongly—even when every candidate shares the same subclass.

Query expansion improves ranking—selectively



0.8193

baseline mAP

0.8417

query expansion mAP

Best setting

$k = 3 \cdot \alpha = 1$
Diffusion peaks at 0.8332.

Neighborhood agreement improves the ranking, but the benefit depends on the query.