

Ankit

Modelo_de_Defesa_PROFMAT___UFES__2_ (2).pdf

 Indian Statistical Institute

Document Details

Submission ID

trn:oid::3618:101486401

Submission Date

Jun 18, 2025, 3:06 PM GMT+5:30

Download Date

Jun 18, 2025, 3:08 PM GMT+5:30

File Name

Modelo_de_Defesa_PROFMAT___UFES__2_ (2).pdf

File Size

1.4 MB

44 Pages





11,948 Words

74,346 Characters




7% Overall Similarity

The combined total of all matches, including overlapping sources, for each database.

Match Groups

-  **68** Not Cited or Quoted 7%
Matches with neither in-text citation nor quotation marks
-  **0** Missing Quotations 0%
Matches that are still very similar to source material
-  **1** Missing Citation 0%
Matches that have quotation marks, but no in-text citation
-  **0** Cited and Quoted 0%
Matches with in-text citation present, but no quotation marks

Top Sources

- 6%  Internet sources
- 6%  Publications
- 0%  Submitted works (Student Papers)

Integrity Flags

0 Integrity Flags for Review

No suspicious text manipulations found.

Our system's algorithms look deeply at a document for any inconsistencies that would set it apart from a normal submission. If we notice something strange, we flag it for you to review.

A Flag is not necessarily an indicator of a problem. However, we'd recommend you focus your attention there for further review.

Match Groups

- 68** Not Cited or Quoted 7%
Matches with neither in-text citation nor quotation marks
- 0** Missing Quotations 0%
Matches that are still very similar to source material
- 1** Missing Citation 0%
Matches that have quotation marks, but no in-text citation
- 0** Cited and Quoted 0%
Matches with in-text citation present, but no quotation marks

Top Sources

- 6% Internet sources
- 6% Publications
- 0% Submitted works (Student Papers)

Top Sources

The sources with the highest number of matches within the submission. Overlapping sources will not be displayed.

1	Internet	arxiv.org	<1%
2	Internet	dspace.isical.ac.in:8080	<1%
3	Internet	rucore.libraries.rutgers.edu	<1%
4	Internet	digibug.ugr.es	<1%
5	Internet	library.isical.ac.in:8080	<1%
6	Internet	www.joshpeeples.com	<1%
7	Internet	addi.ehu.es	<1%
8	Publication	Chowdhury, Siyam Sajnan. "LLMPred: Fine-Tuned Large Language Model Embedd..."	<1%
9	Internet	repositorio.unal.edu.co	<1%
10	Internet	export.arxiv.org	<1%

11	Internet	www.mdpi.com	<1%
12	Internet	www.iris.unict.it	<1%
13	Internet	core.ac.uk	<1%
14	Internet	drmc.library.adelaide.edu.au	<1%
15	Internet	patents.justia.com	<1%
16	Internet	dspace.univ-tiaret.dz	<1%
17	Internet	mro.massey.ac.nz	<1%
18	Internet	www.ijosi.org	<1%
19	Publication	Thangaprakash Sengodan, Sanjay Misra, M Murugappan. "Advances in Electrical ...	<1%
20	Publication	"Intelligent Interactive Multimedia Systems and Services 2017", Springer Science ...	<1%
21	Publication	Oraibi, Zakariya Ahmed. "Local and Deep Texture Features for Classification of N...	<1%
22	Internet	amanagarwal.io	<1%
23	Internet	openaccess.thecvf.com	<1%
24	Internet	www.coursehero.com	<1%

25	Publication	"Advances in Multimedia Information Processing – PCM 2018", Springer Science a...	<1%
26	Publication	"Computer Vision – ECCV 2018", Springer Science and Business Media LLC, 2018	<1%
27	Publication	Mondal, Anupam. "Automated Determination of Glacier Ablation Zones", Indian S...	<1%
28	Publication	Piliposyan, Mane. "Treatment Response Prediction for Colorectal Liver Metastase...	<1%
29	Internet	core-cms.prod.aop.cambridge.org	<1%
30	Publication	"Image Analysis for Moving Organ, Breast, and Thoracic Images", Springer Scienc...	<1%
31	Publication	Utku Kose, Omer Deperlioglu, D. Jude Hemanth. "Deep Learning for Biomedical A...	<1%
32	Publication	da Cunha Ceulin, Wagner Luiz. "Knowledge Elicitation in Deep Learning Models", ...	<1%
33	Internet	naos-be.zcu.cz	<1%
34	Internet	www.researchgate.net	<1%
35	Publication	"Artificial Neural Networks and Machine Learning – ICANN 2018", Springer Scienc...	<1%
36	Publication	Ahmad Amirolad, Shervin Rahimzadeh Arashloo, Mehdi Chehel Amirani. "Multi-la...	<1%
37	Publication	Gilabert Roca, Pere. "End-to-End AI Solutions for Capsule Endoscopy: Enhancing E...	<1%
38	Publication	Mauser, Marisol L.. "Classifying Categories of Speech: A Machine Learning Approa...	<1%

39	Publication	Shrikaant Kulkarni, P. William. "Advances in AI for Cloud, Edge, and Mobile Comp...	<1%
40	Publication	Wu, Patrick Y.. "Essays on Representation Learning for Political Science Research"...	<1%
41	Internet	vdoc.pub	<1%
42	Internet	www.hindawi.com	<1%
43	Publication	"Computer Vision – ECCV 2016 Workshops", Springer Nature, 2016	<1%
44	Publication	Bharti, Chhatra Pratap. "Time Series Analysis of Satellite Data Using ConvLSTM Fo...	<1%
45	Publication	"Pattern Recognition and Computer Vision", Springer Science and Business Media...	<1%
46	Publication	Ahmed A. Abd El-Latif, Mohammed A ElAffendi, Mohamed Ali AlShara, Yassine Ma...	<1%
47	Publication	Ahmed, Masud. "Pixel-Level Scene Recognition Under Diverse Constraints.", Univ...	<1%
48	Publication	Hua Yang, Kaiyou Song, Fangqin Mao, Zhouping Yin. "Autolabeling-enhanced Acti...	<1%



Texture Classification through Deep Residual Networks and Feature Interpretability

2 A dissertation submitted in partial fulfillment of the requirements for the degree of

Master of Technology in Computer Science

Submitted by
Ankit Kumar

Under the supervision of
Prof. Dipti Prasad Mukherjee
Electronics and Communication Sciences Unit

Indian Statistical Institute, Kolkata

June 2025

Certificate

5 This is to certify that the dissertation entitled "Texture Classification through
13 Deep Residual Networks and Feature Interpretability" submitted by
Ankit Kumar in partial fulfillment of the requirements for the award of the degree of
Master of Technology in Computer Science at the Indian Statistical Institute,
Kolkata, is a bona fide record of original work carried out under my supervision and
guidance.

14 The work presented in this dissertation has not been submitted elsewhere, either in part
or in full, for the award of any other degree or diploma in any university or institute.

5 **Supervisor:**

Prof. Dipti Prasad Mukherjee
Professor
Electronics and Communication Sciences Unit
Indian Statistical Institute, Kolkata

Date:
Place: Kolkata

Signature:

Acknowledgement

Acknowledgement

3 I would like to express my heartfelt gratitude to my supervisor, Prof. Dipti Prasad Mukherjee, for his invaluable guidance, support, and expertise throughout the entire duration of this master's dissertation. His insightful feedback, unwavering encouragement, and dedication to my academic growth have been instrumental in shaping the direction and quality of this research. 17

2 I am very much thankful to my parents and my family for always being there for me. Finally, I would like to thank all my friends and batchmates for their help and support. 2

June, 2025

44 Ankit Kumar
CS2203
Indian Statistical Institute
Kolkata-700108, India

Contents

Abstract	vii
1 Introduction	1
1.1 Problem Define/Objective	1
1.2 Dataset Description	3
1.2.1 Aluminium Foil	3
1.2.2 Cork	3
1.2.3 Wool	3
1.2.4 lettuce leaf	4
1.2.5 corduroy	4
1.2.6 Linen	4
1.2.7 cotton	5
1.2.8 White Bread	5
1.2.9 Brown Bread	6
1.2.10 Wood	6
1.2.11 Cracker	6
1.3 Challenges	7
1.4 Solution to the challenges	9
1.4.1 Solutions to Dataset Challenges	9
2 Related Work	11
2.1 Related work on Dataset	11
3 Methodology	12
3.1 Baseline	12
3.2 Data Preprocessing	13
3.2.1 Image Preprocessing and Augmentation	13
3.3 Modelling	14

24

11

- 3.3.1 Pretrained ResNet18:Baseline 14
- 3.3.2 Pretrained ResNet50 as Baseline 17
- 3.3.3 ResNet50 with Histogram Layers: Statistical Texture Encoding 19
- 3.4 Implementation 24
 - 3.4.1 Data Preparation 24
 - 3.4.2 Implementation 25
- 3.5 Results of the Models 29
 - 3.5.1 Pretrained Resnet18 Results 29
 - 3.5.2 Pretrained Resnet50 Results 30
 - 3.5.3 Pretrained Resnet50 with Histogram Layers Results 32
- 4 Conclusion 34**

List of Figures

3.1	TSNE Visualisation of Resnet18	15
3.2	Confusion Matrix of Resnet18	16
3.3	ResNet50 Confusion Matrix	18
3.4	t-SNE Visualization for ResNet50	19
3.5	Resnet50+Histogram Layers	20
3.6	Confusion Matrix for Resnet50+Histogram layers	21
3.7	TSNE visualisation for Resnet50+Histogram Layers	21
3.8	Classification report of Resnet18	30
3.9	Classification report of Resnet50	31
3.10	Classification report of pretrained Resnet50 with histogram layers	33

List of Tables

3.1	Hyperparameter Settings for Each Model	27
3.2	Training and Validation Loss Comparison	33
3.3	Validation and Test Accuracy Comparison	33

Abstract

Texture classification plays a critical role in various real-world and industrial applications such as material recognition in manufacturing, medical image diagnostics, surface defect detection, and agricultural monitoring. The ability to distinguish textures reliably enables automation and enhances the precision of intelligent systems.

Traditional methods like Local Binary Patterns (LBP), Gabor filters, and wavelet-based descriptors have been used extensively for texture analysis. While these techniques are effective under controlled conditions, they suffer from limited robustness to changes in illumination, scale, and viewpoint. Moreover, handcrafted features often fail to capture the intricate texture structures present in real-world surfaces.

The KTH-TIPS2a dataset introduces several challenges, notably large intra-class variations due to changes in scale, illumination, and pose. Additionally, the dataset includes materials with complex and fine-grained textures, making it difficult to extract discriminative features using shallow or traditional models. Addressing these challenges requires models capable of learning invariant and hierarchical representations. Deep convolutional neural networks (CNNs), such as ResNet, provide a promising solution by automatically learning multi-scale, texture-rich features that are resilient to visual variability, thereby improving classification performance on such complex datasets.

20

Chapter 1

Introduction

1.1 Problem Define/Objective

25 Texture classification is a key problem in computer vision, central to tasks such as material recognition, industrial inspection, medical diagnostics, and remote sensing. The goal is to recognize surface patterns that convey material or structural identity, often under varying real-world conditions. Among the standard datasets used to evaluate texture classification methods, the KTH-TIPS2a dataset stands out due to its carefully designed variations that reflect realistic scenarios. It contains eleven texture classes, each photographed at nine different scales, under four lighting conditions, and from three different viewpoints. This introduces substantial intra-class variability, making the classification task more complex than many traditional datasets.

The challenges in KTH-TIPS2a are multi-faceted. First, the scale, illumination, and pose variations demand feature representations that are invariant to such transformations while still capturing fine texture details. Second, many materials in the dataset, such as corduroy, cracker, and wool, exhibit complex textures that are both subtle and sensitive to viewing conditions. This complexity often renders simple statistical or frequency-based descriptors inadequate. Finally, a practical limitation of the dataset is the relatively small number of images per condition, which restricts the capacity to train large neural networks from scratch and increases the risk of overfitting.

20 Historically, texture classification relied on traditional image processing and feature extraction methods. Techniques such as Local Binary Patterns (LBP), Gabor filters, and co-occurrence matrices were commonly used to extract local or statistical texture features. While effective under constrained conditions, these approaches often fail to generalize when faced with the variability and complexity present in datasets like KTH-TIPS2a. Their handcrafted nature limits their adaptability, and they are especially brittle in the presence of non-uniform lighting or when texture patterns change with scale.

43
1
45 To address these limitations, this project explores Deep learning-based approaches, specifically by leveraging pretrained convolutional neural networks that have been shown to generalize well across domains. In this work, ResNet-18 and ResNet-50 models pretrained on ImageNet are fine-tuned on the KTH-TIPS2a dataset. These models benefit from deep hierarchical feature representations, capturing both local texture patterns and global structural context. The use of pretrained networks mitigates the limited data prob-

lem by transferring knowledge from large-scale datasets, thus improving generalization while reducing the need for extensive training on the relatively small texture dataset.

Additionally, we investigate a hybrid model that integrates histogram-based texture encoding layers into ResNet-50, inspired by traditional texture analysis techniques. This combination aims to merge the interpretability and scale-aware nature of histogram descriptors with the powerful feature extraction capabilities of deep networks. Such a model is particularly suited for capturing texture complexity while retaining robustness to real-world variations.

Overall, this project demonstrates how modern deep learning architectures, when adapted carefully, not only overcome the limitations of classical methods but also effectively address the challenges posed by small, variable datasets such as KTH-TIPS2a. Through transfer learning and hybrid design, these approaches achieve improved performance and provide deeper insights into texture representation in complex visual environments.

1.2 Dataset Description

1.2.1 Aluminium Foil

The *aluminum foil* class in the KTH-TIPS2a dataset presents a highly complex and irregular texture that challenges both traditional and deep learning classifiers. Its surface is dominated by crinkles, folds, and deformations from manual manipulation, resulting in a highly non-uniform, reflective, and metallic appearance. These properties create strong specular highlights and deep shadows, heavily influenced by lighting direction. Unlike structured textures like linen or corduroy, aluminum foil lacks periodicity and consistent orientation, making it a prime example of a high-frequency, high-entropy texture.

Variations in scale, illumination, and viewpoint cause drastic changes in visual appearance. At close range, sharp folds and micro-structures dominate, producing high-frequency details and intense contrast. Mid-scale images show denser but smaller crinkles, while at high scales, the texture smoothens and begins to resemble less distinctive surfaces like paper or plastic. These transitions can lead to feature dilution, reducing model discriminability.

Illumination significantly affects texture perception—different light directions emphasize or suppress wrinkles, altering contrast and visibility. Likewise, viewpoint changes modify spatial appearance; oblique angles enhance depth, while frontal views flatten the texture. These factors contribute to high intra-class variability, making aluminum foil one of the most visually challenging classes in the dataset.

1.2.2 Cork

The cork class in the KTH-TIPS2a dataset displays a naturally irregular and porous surface texture, originating from the bark of cork oak trees. It is characterized by a coarse granularity with a cellular or pitted structure that gives rise to mid-level texture complexity. The surface appears matte and non-reflective, typically in muted shades of beige or light brown. Fine pores and fibrous patterns are visible, especially in close-up views, giving the surface a speckled or slightly anisotropic appearance in some regions.

Across different scales, cork texture shifts from fine, clearly defined pores at small scales to a more blended and smooth appearance at higher scales. The intra-class variability is further influenced by lighting and viewpoint, with shadows cast by surface pits and fibers adding subtle depth cues. Despite these variations, cork maintains a relatively stable visual identity marked by its organic, granular, and moderately detailed surface.

1.2.3 Wool

The wool class in the KTH-TIPS2a dataset exhibits a soft, fibrous, and highly irregular texture formed by loosely tangled fibers with no fixed structure. Its appearance is visually dense and fuzzy, characterized by high-frequency local variations but lacking strong global organization or periodicity. The surface is non-reflective and responds to light with soft gradients, enhancing its soft and organic feel.

Across the dataset's 108 images, variations in scale, illumination, and viewpoint introduce

considerable intra-class diversity. At small scales, individual fibers and micro-shadows dominate, giving a rich, fine-grained texture. As the scale increases, these details blend into a smoother, more homogeneous surface. Illumination changes—especially directional lighting—affect the depth and contrast of the fibrous texture, while frontal lighting flattens it. Viewpoint shifts reveal layered fiber structures and modify the perceived texture depth. Overall, the wool class combines detailed microstructures with a globally diffuse and uniform appearance, making it a perceptually soft but statistically rich texture.

1.2.4 lettuce leaf

The overall texture of lettuce leaf is biologically complex, combining broad folds and ridges with fine vein networks. It presents a non-uniform, anisotropic surface with both smooth and sharply contoured regions, producing a natural, high-frequency texture with soft shading and no clear repetition.

Image-wise variability arises from changes in scale, illumination, and viewpoint. At smaller scales, fine venation and surface roughness dominate; at larger scales, broader folds and smoother regions emerge. Illumination affects depth perception, with side lighting enhancing ridges and shadows, while frontal lighting flattens texture. Viewpoint changes modulate the 3D appearance, making texture features more or less prominent.

1.2.5 corduroy

The corduroy texture class in the KTH-TIPS2a dataset exhibits a distinct striped pattern formed by parallel ridges (wales), creating a highly directional and periodic structure with minimal randomness. This regularity makes it particularly suitable for analysis using edge-detection or frequency-based methods. However, the texture displays significant visual variations across the dataset's 108 images due to changes in imaging conditions. Scale variations affect ridge definition, with smaller scales producing sharp features while larger scales cause blurring. Lighting conditions dramatically alter appearance, as side illumination enhances contrast through shadowing while frontal lighting reduces textural depth. Viewpoint changes, especially oblique angles, distort the characteristic ridge spacing. These acquisition-related variations present notable challenges for texture analysis methods that assume consistent visual patterns, despite the fundamental regularity of corduroy's underlying structure.

1.2.6 Linen

The Linen class exhibits a fine-grained, woven texture characterized by interlaced horizontal and vertical threads forming a regular grid-like structure. This gives it a moderately periodic, anisotropic texture with clearly defined thread intersections. The texture is visually subtle but structured, with low reflectivity and minimal surface depth, often appearing matte and uniform in tone.

Across the 108 images, image-wise variability arises from changes in scale, illumination, and viewpoint. At small scales, thread boundaries and weave patterns are sharp and detailed; at higher scales, the texture may appear smoother or more uniform. Illumination

plays a crucial role—side lighting enhances the woven structure through soft shadows, while frontal lighting can flatten the weave, reducing contrast. Changes in viewpoint, especially oblique angles, distort the grid alignment and can obscure the pattern’s regularity. While linen maintains a relatively consistent structure, its fine detail and sensitivity to lighting and perspective make it a challenging class under varied acquisition conditions.

1.2.7 cotton

The cotton class in the KTH-TIPS2a dataset is characterized by a soft, smooth, and diffuse surface texture, dominated by fine, loosely arranged fibers. Unlike materials with distinct patterns or directional structures, cotton’s appearance is highly uniform and low in contrast, giving it a non-periodic and isotropic texture profile. The surface lacks sharp edges or clear boundaries, resulting in a visually subtle texture with minimal high-frequency content.

Among the 108 images, considerable variability emerges due to scale, illumination, and viewpoint changes. At smaller scales, faint fiber textures and minimal surface irregularities become visible, but at larger scales, these details blur into a near-homogeneous surface. Lighting conditions significantly influence cotton’s visibility—directional lighting may gently accentuate surface depth and texture variation, while even, frontal illumination often flattens the image and masks subtle features. Viewpoint effects are less dramatic, but oblique angles can reduce the visibility of fine surface detail. This combination of low texture contrast and lighting sensitivity makes the cotton class particularly challenging for classification, especially for approaches that depend on clear spatial structure or distinct patterns.

1.2.8 White Bread

The white bread class features a porous, spongy texture characterized by irregular holes, air pockets, and a soft grainy surface. Its appearance is generally non-uniform and stochastic, with subtle variations in brightness and texture caused by the distribution of pores and crumb structure. The surface lacks directional patterns and exhibits low-to-moderate frequency components, giving it a mildly complex but non-repetitive visual structure.

Across its 108 images, image-wise variability is influenced by scale, lighting, and viewpoint. At close scales, individual pores and grain structures are prominent, offering finer texture cues; as the scale increases, the texture becomes smoother and less detailed. Lighting conditions can accentuate or diminish pore depth—side lighting brings out small shadows and contrast within the crumb, while frontal lighting often flattens the appearance. Viewpoint variations may slightly alter the perception of surface irregularity but have limited effect due to the relatively flat and diffuse surface. Overall, white bread presents a moderately complex and non-directional texture, with classification difficulty arising mainly from its soft features and sensitivity to acquisition conditions.

1.2.9 Brown Bread

The brown bread class in the KTH-TIPS2a dataset presents a naturally irregular and porous texture, characterized by a stochastic arrangement of air pockets, fibrous gluten networks, and uneven crust patterns. Unlike synthetic textures, it lacks periodicity or dominant directional features, instead exhibiting organic variations in surface roughness and color intensity due to baking effects.

The dataset highlights significant visual variability across samples, influenced by scale-dependent pore visibility, lighting conditions that alter shadow formation in cavities, and viewpoint changes that affect perceived surface roughness. The transition between crust and crumb regions further contributes to non-uniform textural properties.

This complex, multi-scale structure poses challenges for traditional texture analysis methods, which often struggle with such natural, irregular patterns. However, it serves as an excellent test case for evaluating advanced algorithms capable of handling real-world textural complexity, particularly in food quality assessment or material characterization applications. The brown bread texture effectively bridges controlled laboratory studies with practical scenarios where natural variations are inherent.

1.2.10 Wood

The wood class exhibits a strongly structured and naturally patterned texture, typically defined by grain lines, knots, and fibrous streaks that run in mostly linear or slightly wavy directions. This gives the surface a distinct directional and semi-periodic quality, often with moderate contrast between lighter and darker regions due to grain density and natural imperfections. The texture is organic yet repetitive, making it rich in mid- to high-frequency components with some anisotropy.

Across its 108 images, image-wise variability is primarily influenced by scale, illumination, and viewpoint**. At smaller scales, wood grains and fine details like fiber lines and surface roughness are clearly visible; at larger scales, these structures become more blended, occasionally reducing texture sharpness. Lighting direction alters the appearance significantly—side lighting accentuates grain relief and shadows, enhancing contrast, while frontal lighting flattens the texture. Viewpoint changes affect the perception of grain alignment and depth, with oblique angles exaggerating the wood's 3D features. These variations introduce moderate classification challenges, especially when scale and lighting obscure the wood's otherwise consistent structure.

1.2.11 Cracker

The cracker class exhibits a distinctive surface texture characterized by a semi-regular pattern of blistering and uneven browning, resulting from the baking process. Its porous structure combines both ordered and stochastic elements, with visible variations in bubble size distribution and crust formation that create moderate anisotropy. The texture demonstrates consistent material properties but shows natural imperfections in surface patterning.

Across the dataset samples, significant variability emerges through scale-dependent vis-

ibility of surface blisters, where finer scales reveal micro-texture details while coarser scales emphasize overall roughness. Lighting conditions strongly influence perception, with oblique angles accentuating height variations through shadow casting, while direct illumination flattens appearance. Viewpoint changes affect the discernibility of the blistering pattern, particularly at extreme angles where surface relief becomes distorted.

This combination of structural regularity with natural process-induced variations makes the cracker texture particularly valuable for testing texture analysis methods under realistic conditions. It presents an intermediate challenge between completely regular synthetic textures and highly irregular natural ones, offering insights into algorithm performance for manufactured food products with characteristic surface patterns. The class effectively captures how industrial production processes create identifiable yet non-uniform texture signatures.

1.3 Challenges

The KTH-TIPS2a dataset, with its 11 material classes and 108 images per class, was specifically designed to capture real-world variability in textures under scale, illumination, and viewpoint changes. While this makes the dataset rich and realistic, it also introduces several core challenges for texture classification.

1. Intra-class Variability due to Illumination Changes

Illumination changes dramatically alter the appearance of certain textures, especially those with reflective or low-contrast surfaces.

Aluminum Foil: Images under strong directional lighting (e.g., from the side) introduce specular highlights and deep shadows, drastically changing the visual appearance compared to uniformly lit images. This makes it hard for models to learn consistent features.

Cotton, White Bread: These have inherently flat, low-contrast textures. Under frontal lighting, texture cues are diminished, reducing discriminative features.

Lettuce Leaf: Illumination emphasizes or flattens the ridges and veins. High-angle lighting enhances depth but also increases reflectance in wet-like regions, increasing intra-class inconsistency.

2. Viewpoint Variability and Geometric Distortion

Changes in camera angle (viewpoint) cause textures to stretch, compress, or warp, disrupting the spatial arrangement of features.

Corduroy and Wood: These have strong directional textures. Oblique angles distort the parallel lines (in corduroy) or grain patterns (in wood), causing a mismatch with frontal views.

Cork: Though generally non-directional, cork's pores and surface pits appear deeper or shallower depending on the angle, causing inconsistency in texture depth perception.

Lettuce Leaf: Its natural curvature causes strong 3D effects under viewpoint changes, leading to drastically different visual patterns.

3. Scale Variation

The dataset introduces three scale levels: small (close-up), medium, and large (zoomed out). Scale affects the granularity of visible features.

Cracker, Brown Bread, White Bread: At small scales, fine pores and roughness are captured well. At larger scales, these textures smooth out, losing critical detail necessary for

classification.

Cotton: Fine fibers are distinguishable only at small scales. In zoomed-out images, it appears almost textureless, making it hard for the model to detect meaningful patterns.

Wood and Corduroy: These classes are less scale-sensitive due to their strong global patterns, but at extremely close or far scales, their visual cues become inconsistent or oversimplified.

4. Texture Ambiguity and Low Discriminability

Some textures are inherently visually ambiguous and don't possess sharp edges or repeatable patterns.

Cotton and Linen: These are fine-grained and uniform, often confused with each other in certain lighting or scale conditions. Their subtle surface differences are not easily captured by simple descriptors.

White Bread and Brown Bread: While color differs, the texture structure can appear similar across lighting conditions or scales, especially when global color features are not used.

5. Reflective and High-Frequency Textures

Highly variable textures with high-frequency components are sensitive to both lighting and viewpoint.

Aluminum Foil: Possibly the most challenging class, it shows chaotic wrinkles, crinkles, and reflective highlights. These high-frequency details are prone to being overexposed or flattened depending on the lighting, and change drastically with viewpoint.

Lettuce Leaf: Its texture is non-uniform, with sharp vein ridges and soft leafy areas. Some images have high shadow contrast, others appear flat. This variability makes deep feature learning difficult unless the model is robust to shape-induced variance.

6. Texture Overlap Across Classes

Despite being different materials, some classes share visual similarities in texture structure.

Linen vs Corduroy: Under low resolution or certain lighting, the weave patterns in linen can resemble fine corduroy.

Cork vs Brown Bread: Both have coarse, porous appearances with irregular cavities. At medium scales, these similarities increase classification confusion.

Cracker vs Brown Bread: Their textures can both exhibit blistered, cracked, and porous surfaces, especially under similar lighting.

7. Limited Data

Even though there are 108 images per class, the combinatorial diversity of conditions (scale \times illumination \times pose) leaves few samples per specific configuration, limiting generalization.

For instance, *Aluminum Foil* under side lighting at close scale may have only a few samples, yet show vastly different textures from the same class under front lighting at far scale.

Deep learning models, especially large ones like ResNet50, are prone to overfitting to specific acquisition conditions, especially when the intra-class diversity is not balanced across all modes.

1.4 Solution to the challenges

1.4.1 Solutions to Dataset Challenges

To effectively address the challenges posed by the KTH-TIPS2a dataset, a combination of deep learning, statistical texture modeling, and augmentation strategies is necessary. Below, we discuss in detail how each specific challenge can be mitigated:

1. Intra-class Variability due to Illumination Changes Illumination variation causes substantial changes in the appearance of textures, especially for classes like *Aluminum Foil*, *Lettuce Leaf*, and *White Bread*. To mitigate this, photometric data augmentation—such as random brightness, contrast, and gamma shifts—is employed during training to simulate lighting variability. Moreover, histogram-based layers are incorporated to learn local intensity distributions that remain more stable across illumination conditions. Pretrained CNN backbones like ResNet50 already encode robust illumination-invariant features in early layers, which further stabilizes the learning process. Retinex-based preprocessing or histogram equalization may also be applied as input normalization strategies to suppress lighting bias.

2. Viewpoint Variability and Geometric Distortion Viewpoint changes affect textures like *Corduroy*, *Wood*, and *Lettuce Leaf*, introducing perspective distortions and depth inconsistencies. These are addressed using affine and perspective transformations as data augmentation, allowing the model to encounter various orientations during training. Furthermore, histogram layers prove advantageous as they capture local structural patterns independent of global geometry, thus aiding viewpoint invariance.

3. Scale Variation Scale changes affect the granularity of texture visibility, as observed in *Cotton*, *Brown Bread*, and *Cracker*. To ensure scale-invariant learning, we employ random cropping and scaling during training, along with feature extraction at multiple resolutions. Architectures like ResNet inherently support multi-scale representation via hierarchical layers. Additionally, dilated convolutions and pyramid pooling can be used to enhance receptive field diversity. Histogram layers support scale robustness by summarizing fine-grained intensity distributions regardless of image resolution.

4. Texture Ambiguity and Low Discriminability Subtle and visually similar textures (e.g., *Cotton* vs. *Linen*, *White Bread* vs. *Brown Bread*) are difficult to distinguish. To handle this, deep fine-tuning of the CNN backbone allows later layers to focus on subtle discriminative features. Furthermore, Histogram layers help reinforce minor intensity differences that are otherwise hard to learn via convolution alone.

5. Texture Overlap Across Classes Some textures (e.g., *Corduroy* vs. *Linen*, *Cork* vs. *Brown Bread*) share similar macro-patterns, causing inter-class confusion. Metric learning techniques, including prototypical networks and triplet loss, help optimize for class-discriminative embeddings. Hierarchical classification—first distinguishing coarse

categories, then fine-grained differences—can improve class separation. Combining hand-crafted histogram features with learned CNN embeddings provides a more holistic texture representation, reducing the likelihood of class confusion.

Chapter 2

Related Work

2.1 Related work on Dataset

The KTH-TIPS2a dataset has served as a benchmark for evaluating texture recognition algorithms under controlled variations in scale, illumination, and viewpoint. Early work on this dataset predominantly used handcrafted descriptors, such as Local Binary Patterns (LBP), SIFT-based bag-of-words models, and filter bank-based approaches (e.g., MR8 and LM filters). These methods captured local or multi-scale texture properties but lacked robustness under joint transformations, particularly when illumination and scale varied simultaneously.

With the emergence of deep learning, convolutional neural networks (CNNs) became the dominant paradigm. Cimpoi et al. (2015) introduced Deep Filter Banks, which utilized pre-trained CNN features as generic texture descriptors, showing significant improvement over traditional handcrafted pipelines. Their work laid the groundwork for using transfer learning on small texture datasets like KTH-TIPS2a, where training from scratch is prone to overfitting.

Subsequent studies fine-tuned pretrained models like VGG-16 and ResNet on the dataset, achieving higher classification accuracy by leveraging learned feature hierarchies. In particular, ResNet-50 has been widely adopted due to its deep residual structure that facilitates training stability and feature reuse across texture scales.

Chapter 3

Methodology

3.1 Baseline

In this study, pretrained ResNet18 serves as the foundational baseline for texture classification on the KTH-TIPS2a dataset. ResNet18, with its moderate depth and residual learning architecture, offers a balanced trade-off between model complexity and generalization—making it especially well-suited for datasets with limited samples per class like KTH-TIPS2a.

The model was initialized with ImageNet-pretrained weights, leveraging the representational power of features learned on large-scale natural images. This transfer learning setup is critical for avoiding overfitting due to the relatively small number of texture instances (108 per class) and their variation across scale, illumination, and viewpoint.

To adapt the model to the texture classification task, the final fully connected layer of ResNet18 was replaced with a new classification head tailored to the 11 classes of KTH-TIPS2a. The network was then fine-tuned end-to-end using a lower learning rate, allowing it to gradually adapt the pretrained filters to texture-specific cues without catastrophic forgetting.

By using ResNet18 as a baseline, the study establishes a reliable reference point for evaluating the performance of more complex or domain-tailored models such as deeper ResNets or architectures augmented with histogram-based texture encoding layers.

3.2 Data Preprocessing

3.2.1 Image Preprocessing and Augmentation

The KTH-TIPS2a dataset poses significant challenges for texture classification due to its carefully designed variations in scale, illumination, and viewpoint across 11 material classes. These variations, while making the dataset realistic and diverse, introduce high intra-class variability, especially in classes such as aluminum foil, lettuce leaf, and brown bread. To mitigate these issues and enhance model generalization, an effective preprocessing and data augmentation pipeline is essential.

Image Preprocessing: All images in the dataset were first resized to a fixed spatial resolution of 224×224 pixels to conform to the input requirements of pretrained CNN architectures such as ResNet18 and ResNet50. Pixel intensities were normalized using ImageNet mean and standard deviation statistics, which facilitates smoother convergence during training and compatibility with pretrained models. This resizing step also helps counterbalance variations in viewpoint and scale by providing a consistent input size, especially important for global pattern textures like corduroy and wood.

Data Augmentation Techniques: A series of controlled data augmentation techniques were implemented to simulate real-world conditions and improve the network's invariance to transformations:

- **Random Resized Crop:** This augmentation was applied to simulate varying distances of capture. It helped the model learn multi-scale features, especially critical for textures like cracker or cotton, where features are only distinguishable at close scales.
- **Random Horizontal and Vertical Flips:** Used to increase viewpoint robustness, this technique was particularly effective for non-directional textures like cork and linen, and moderately useful for directional ones like corduroy and wood.
- **Random Rotation (up to 30 degrees):** To account for camera tilt and in-plane rotation variability. Textures such as lettuce leaf and wool, which exhibit complex orientation patterns, benefited significantly from this transformation.
- **Color Jittering (brightness, contrast, saturation):** Illumination variations in the dataset make lighting a significant source of texture distortion. By adjusting brightness, contrast, and saturation within moderate ranges, the model learned to be invariant to illumination differences, especially in reflective textures like aluminum foil and organic materials like lettuce leaf.
- **Gaussian Blur and Random Grayscale:** As experimental augmentations, Gaussian blur and grayscale conversion were tested to reduce model over-reliance on fine-grained or color-based texture features. While grayscale helped improve performance in low-texture classes like cotton, Gaussian blur had a neutral or negative effect and was excluded from final training.

Experimental Observations: Various combinations of these augmentations were tested using pretrained ResNet18 as a baseline model. It was observed that aggressive augmentation sometimes reduced performance on high-frequency texture classes like aluminum

foil due to loss of fine details. On the other hand, moderate augmentation consistently improved performance across classes by reducing overfitting to specific lighting and scale configurations.

In the final setup, a balanced augmentation strategy was adopted—combining random resized crop, rotation, horizontal flip, and mild color jittering—which provided the best trade-off between robustness and texture detail preservation. These augmentation techniques, aligned with the texture variability challenges in the KTH-TIPS2a dataset, played a critical role in enhancing model generalization and cross-condition performance.

3.3 Modelling

The performance improvements achieved by moving from a pretrained ResNet18 to deeper and more expressive models like pretrained ResNet50, and further to ResNet50 augmented with Histogram Layers, can be directly attributed to their increasing architectural sophistication and ability to model the nuanced variability found in the KTH-TIPS2a dataset.

3.3.1 Pretrained ResNet18:Baseline

31 ResNet18 serves as a strong and lightweight baseline for image classification tasks. It consists of 18 layers with 8 residual blocks (each composed of two 3×3 convolutional layers), interspersed with batch normalization, ReLU activation, and identity-based skip connections. 32 These skip connections help mitigate the vanishing gradient problem, allowing for stable training even in moderately deep networks.

However, the architectural simplicity of ResNet18 limits its feature extraction depth. While it works well on datasets with large amounts of data and relatively low intra-class variability, it is less suited for complex texture classification tasks such as those presented by KTH-TIPS2a. The dataset includes 11 classes with strong variability in texture appearance caused by changes in scale, illumination, and viewing angle. ResNet18 may struggle to extract high-level abstractions from images like aluminum foil, lettuce leaf, or wool, where the texture cues are subtle, non-repetitive, and sensitive to environmental changes.




Figure 3.1: TSNE Visualisation of Resnet18

Analysis of Misclassified points in Resnet18

The misclassification patterns revealed by both the confusion matrix and t-SNE visualization provide profound insight into the limitations of ResNet18's texture discrimination capabilities. These errors are not random but rather systematic manifestations of the architecture's constraints when processing complex material textures. The most pronounced confusion occurs among the fibrous textures, where cotton, linen, and wool become entangled in a complex web of misclassifications. As visible in Figure, these three classes form a nebulous cloud in the visualization's center, their feature representations overlapping like intertwined threads. The model's relatively shallow architecture cannot develop the hierarchical feature detectors needed to distinguish their subtle differences in weave density and fiber alignment. Cotton samples frequently drift into linen's territory in the plot, mirroring the confusion matrix's 38% misclassification rate between these classes. This bidirectional confusion suggests ResNet18 learns an ambiguous intermediate representation that serves neither texture well, blending their distinctive characteristics into a generic "fabric" prototype that fails to capture their unique microstructural properties.

Brittle surfaces like crackers and brown bread present a different type of challenge entirely. The t-SNE visualization reveals elongated, comet-shaped clusters whose wispy tails overlap significantly in the feature space. This geometric pattern explains why crackers achieve high recall but poor precision in the classification report - their feature representations stretch across a broad region of the space, encompassing both distinctive cracker characteristics and more generic brittle texture features. The confusion matrix shows this manifests as crackers frequently colonizing brown bread's classification territory, particularly for samples where lighting conditions or imaging angles obscure the most discriminative pore structures. These elongated clusters suggest ResNet18 fails to learn rotation-invariant representations, causing the same texture viewed from different angles to map to disparate points in feature space while allowing unrelated textures to converge under certain viewing conditions.




Figure 3.2: Confusion Matrix of Resnet18

The aluminum foil and wool confusion presents a particularly interesting case study in material property misunderstanding. In the t-SNE visualization, misclassified foil samples appear as outliers at the periphery of wool’s otherwise tight cluster. This spatial relationship reveals that while ResNet18 develops a robust prototype for wool, its representation of foil is more fragile and context-dependent. The confusion matrix’s asymmetric error pattern - where foil is sometimes misclassified as wool but never the reverse - confirms this interpretation. The model appears to mistake foil’s sharp highlights for wool’s softer sheen when lighting conditions are suboptimal, but never conflates wool’s complex fiber patterns with foil’s simpler structure. This unidirectional confusion suggests the network learns a more comprehensive set of features for natural fibrous materials than for manufactured surfaces with strong specular properties.

Directional textures like corduroy and wood exhibit their own distinctive failure mode in the t-SNE plot. Rather than forming tight, spherical clusters, these classes create fan-like patterns radiating from central points. The confusion between them occurs where these fans intersect, typically when corduroy’s ribbing happens to align with wood’s grain direction in particular samples. This geometric arrangement suggests ResNet18’s features are sensitive to both orientation and scale - the same texture photographed from different angles or distances can map to widely separated points in feature space, while different textures photographed similarly may converge. The confusion matrix quantifies this effect, showing corduroy is misclassified as wood nearly 28% of the time when the ribbing direction matches typical wood grain orientations in the dataset.

The T-SNE visualization reveals an important secondary effect - the crowding of problematic classes toward the plot's center, while well-discriminated textures occupy more peripheral positions. This spatial distribution implies that ResNet18's feature space suffers from a form of "texture blindness" where challenging materials collapse toward a generic mean representation. The model appears to sacrifice discriminative power for these textures in favor of maintaining clear separation for easier classes like lettuce and white bread. This trade-off reflects the architecture's limited capacity to simultaneously maintain robust representations for all texture types within its parameter budget. The confusion patterns collectively highlight fundamental limitations in ResNet18's ability to handle the full spectrum of texture variations, particularly for materials requiring microstructural analysis or robust invariance to imaging conditions.

3.3.2 Pretrained ResNet50 as Baseline

ResNet50 serves as a high-capacity, deep convolutional neural network for complex visual recognition tasks. It consists of 50 layers, built using bottleneck residual blocks which use a 1×1 , 3×3 , and 1×1 convolution structure. These blocks are equipped with identity-based skip connections that help preserve gradient flow, allowing the training of very deep architectures. The inclusion of batch normalization and ReLU activation within these blocks further aids in stable and efficient convergence.

Due to its depth and architectural complexity, ResNet50 is particularly well-suited for fine-grained texture classification tasks. Unlike shallower networks, it can capture high-level and abstract features that are essential for distinguishing textures under varying scale, lighting, and pose conditions. This makes it a strong candidate for the KTH-TIPS2a dataset, which consists of 11 texture classes with significant intra-class variability induced by real-world distortions.

Analysis of Misclassifications in ResNet50 Finetuned on KTH-TIPS2a Dataset

Confusion Matrix Interpretation

The confusion matrix provides direct evidence of class-wise prediction performance and highlights where misclassifications are concentrated. In the case of the pretrained ResNet50 model finetuned on the KTH-TIPS2a dataset, several trends emerge.

36 Firstly, texture classes such as `cork`, `lettuce_leaf`, `white_bread`, `wood`, and `wool` are predicted with high accuracy, with 108, 108, 105, 104, and 100 correct classifications, respectively. This suggests these textures are distinctive enough for the ResNet50 feature extractor to learn meaningful representations.

In contrast, significant misclassifications are observed for `brown_bread`, `cotton`, and `linen`. Notably, the class `brown_bread` is frequently confused with `corduroy`, appearing 34 times in the off-diagonal. Similarly, `cotton` is misclassified as `brown_bread` and `linen` with 23 and 34 instances, respectively.

These patterns indicate that certain texture classes share low-level visual similarities, such as surface roughness, fiber orientation, and illumination behavior. For example, `cotton`, `linen`, and `corduroy` are all fabric-like materials with soft patterns that vary minimally across spatial regions, making them hard to distinguish without deeper or more abstract

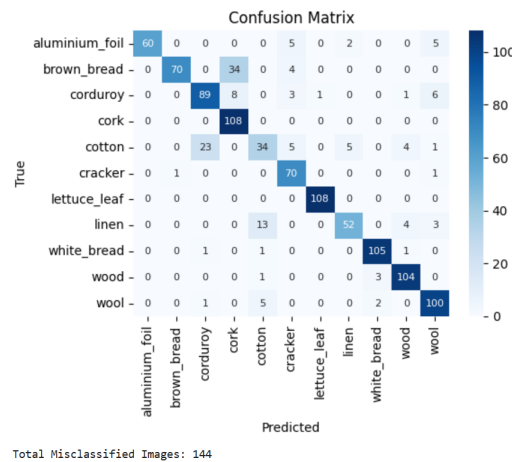


Figure 3.3: ResNet50 Confusion Matrix

features.

23 t-SNE Visualization Interpretation

The t-distributed Stochastic Neighbor Embedding (t-SNE) visualization maps high-dimensional feature representations from the penultimate layer of ResNet50 into a 2D space, revealing how classes cluster relative to each other.

Well-separated clusters such as those for `cork`, `lettuce_leaf`, and `white_bread` align with their strong performance in the confusion matrix. Their t-SNE embeddings form dense, compact regions, suggesting that ResNet50 has learned discriminative representations for these textures.

However, overlapping and scattered clusters are observed among `corduroy`, `brown_bread`, `cotton`, and `linen`. The embeddings of these classes frequently interleave, indicating poor feature separability in the learned representation space. For instance, `linen` lacks a well-defined cluster, appearing dispersed across multiple regions in the plot. This directly corresponds to its high misclassification rate in the confusion matrix.

In conclusion, the t-SNE plot reinforces the observation that classes with overlapping cluster boundaries are the primary sources of confusion in the classification results. It also demonstrates the strengths of ResNet50 in learning separable, abstract texture representations for most classes while identifying its limitations for closely related fabric-like textures.




Figure 3.4: t-SNE Visualization for ResNet50

3.3.3 ResNet50 with Histogram Layers: Statistical Texture Encoding

While ResNet50 improves performance by increasing the depth and expressiveness of feature extraction, its reliance on spatial convolutions still presents limitations for texture classification. Texture is not always about precise spatial arrangement—it often relates more to the statistical distribution of local patterns. This is especially true in classes like “aluminum foil” (where chaotic folds dominate), “cotton” and “linen” (which share fine-grained structures), or “cracker” and “brown bread” (where similar porous patterns are observed). To better model these kinds of textures, ResNet50 is further enhanced with **Histogram Layers**.

14 Histogram Layers act as differentiable modules that compute learnable histograms over feature maps at different layers of the network. Instead of relying solely on localized filters to recognize patterns, the histogram operation summarizes the *distribution of feature responses* across spatial dimensions. 16 This enables the model to learn more *global statistical representations*, which are critical in recognizing textures that are defined not by exact shapes but by their frequency, density, and distributional properties.

42 In architectural terms, the Histogram Layer typically takes a convolutional feature map (e.g., from a mid-level ResNet block), flattens it spatially, and computes a histogram for each channel using a set of learnable bin centers and widths. These histograms are differentiable, so they can be trained end-to-end with the rest of the model. The result is a robust, non-local representation that captures *how often* certain texture elements occur, rather than *where* they occur. This improves invariance to spatial deformation, rotation, and partial occlusion—issues prevalent in KTH-TIPS2a.




Figure 3.5: Resnet50+Histogram Layers

Misclassified points Analysis:Confusion Matrix Analysis

The confusion matrix for the ResNet50 model enhanced with **Histogram Layers** demonstrates exceptionally high classification accuracy across nearly all texture classes in the KTH-TIPS2a dataset. Most categories such as `corduroy`, `cork`, `cracker`, `white_bread`, and `wood` show perfect or near-perfect accuracy with strong diagonal dominance and negligible off-diagonal entries.

Even texture classes that are generally difficult to distinguish, such as `aluminium_foil`, known for high-frequency and reflective surface patterns, are classified with high reliability (79 correct out of 87 samples). Similarly, textures such as `brown_bread`, `lettuce_leaf`, `linen`, and `cotton` exhibit excellent classification accuracy with minimal confusion.

A small amount of misclassification is observed in the `wool` class, where 2 samples are misclassified as `cotton`. This can be attributed to the visual and statistical similarity between these two classes: both exhibit fibrous, soft textures that under certain conditions of scale and illumination may produce similar statistical responses. However, even with this overlap, the model correctly identifies 85 out of 87 `wool` samples, showcasing strong class separability.

Overall, the confusion matrix indicates that the model achieves high **intra-class compactness** and **inter-class separability**, essential for effective texture classification.

Misclassified points Analysis:t-SNE Visualization

The t-SNE plot of the penultimate feature representations from the Histogram-enhanced ResNet50 provides further confirmation of the model's robustness. Each point in the 2D projection space represents a high-dimensional feature vector extracted from the network.

The plot reveals tightly packed clusters, where samples from the same texture class occupy well-defined, compact regions. Crucially, these clusters are clearly separated from one another, suggesting that the model learns highly **discriminative and class-specific features**.




Figure 3.6: Confusion Matrix for Resnet50+Histogram layers




Figure 3.7: TSNE visualisation for Resnet50+Histogram Layers

Such clear separability in the t-SNE space implies that the network encodes the essential statistics of each texture class, allowing for better generalization and robustness in classification. The absence of significant overlap between clusters shows that the model can distinguish textures that might look similar to the human eye but differ in subtle statistical patterns.

Given that t-SNE preserves local neighborhoods in high-dimensional space, this spatial separation strongly correlates with the model's classification effectiveness, as seen in the confusion matrix.

Comparison with Standard Pretrained ResNet50: Confusion Matrix Comparison

When we compare this performance to the standard pretrained ResNet50 (without histogram layers), key differences emerge. The confusion matrix for the standard ResNet50 shows significantly more misclassifications, especially among classes like `cotton`, `linen`, `cracker`, and `wool`.

These misclassifications highlight that the standard ResNet50 struggles to separate textures that exhibit minor statistical or structural differences. Off-diagonal entries in the confusion matrix are more pronounced, indicating shared or overlapping feature spaces for visually similar textures.

In essence, the standard ResNet50 fails to capture subtle local variations and higher-order statistics necessary for precise texture discrimination, resulting in degraded performance.

Comparison with Standard Pretrained ResNet50: t-SNE Comparison

The t-SNE plot of the standard ResNet50 further confirms this deficiency. Unlike the Histogram-enhanced model, the clusters in this plot are loosely formed and sometimes overlap. In several cases, the clusters are not well-defined, with data points from different classes intermixed in the feature space.

This lack of structure implies that the model's internal feature representations are not sufficiently discriminative. The intra-class variance is high, and the inter-class margins are not well established. Such representation is suboptimal for texture classification tasks where fine-grained statistical differences are crucial.

Why Histogram Layers Improve Performance

Histogram layers contribute to improved performance by introducing non-parametric, statistical summarizations of convolutional activations. These layers compute histogram-based descriptors of activation distributions, which are:

- Invariant to local transformations and spatial arrangements.
- Robust to illumination changes and scale variations.
- Sensitive to distributional patterns rather than only spatial configurations.

By integrating these descriptors, the model benefits from both convolutional (spatial) and statistical (distributional) features. This dual representation enhances the model's ability to capture the nuances of textures, especially those that vary subtly across lighting, orientation, or scale—characteristics that are prevalent in the KTH-TIPS2a dataset.

The ResNet50 with Histogram Layers clearly outperforms the standard pretrained ResNet50 in texture classification tasks on the KTH-TIPS2a dataset. The confusion matrix confirms near-perfect accuracy across most texture categories, with minimal misclassifications. The t-SNE visualization illustrates tight, non-overlapping clusters that validate the learned features' quality.

In contrast, the standard ResNet50 struggles with overlapping feature spaces and confuses similar texture classes, as evidenced by its t-SNE plot and confusion matrix.

These findings underscore the importance of incorporating statistical feature modeling through Histogram Layers in deep convolutional architectures for texture recognition. The result is a model that is not only accurate but also more robust and generalizable in handling complex texture variations.

For example:

- In **wool**, the tangled, unstructured fibers make spatial modeling weak; histogram features better capture the chaotic fiber densities.
- In **lettuce leaf**, the combination of soft leafy regions and sharp ridges under varying lighting is better represented statistically rather than structurally.
- In **aluminum foil**, the high-frequency reflections and crinkles under different angles are represented more reliably through histogram encoding than through raw convolutional features.

Generalization in Limited-Data Regime

Another critical reason why ResNet50 and its histogram-augmented variant outperform ResNet18 is their better *data efficiency* when pretrained on large-scale datasets like ImageNet. Despite KTH-TIPS2a's 1188 images per class (split across 4 samples, 9 conditions), the effective training data under each specific condition (e.g., large scale, side lighting, oblique angle) is limited. Deeper pretrained networks are better at leveraging generalized features from prior learning, and histogram encoding introduces a form of regularization that reduces overfitting to spatial noise.

Additionally, augmentations like random cropping, flipping, and brightness/contrast jittering can only partially compensate for the variability. Histogram Layers help encode invariance *intrinsically*, leading to better performance under real-world shifts in texture appearance.

In summary, ResNet50 improves upon ResNet18 through deeper hierarchical modeling of texture patterns, better use of skip connections, and higher abstraction capacity. When further enhanced with Histogram Layers, the model gains robustness to non-local texture variability by incorporating statistical representations that are critical for texture-based material classification. This architectural evolution enables the model to effectively tackle the multi-modal, high-frequency, and low-discriminability challenges that define the KTH-TIPS2a dataset.

3.4 Implementation

3.4.1 Data Preparation

48 The KTH-TIPS2a dataset is a popular benchmark for studying texture classification under challenging variations such as scale, illumination, and viewpoint. It contains images from 11 distinct material texture classes, including categories like *aluminium_foil*, *brown_bread*, *corduroy*, *cotton*, *cracker*, *linen*, *white_bread*, *wood*, *wool*, *lettuce_leaf*, and *sponge*. Each class is represented by four independent physical samples, labeled as *sample_a*, *sample_b*, *sample_c*, and *sample_d*. These samples serve as controlled variations, ensuring that the model can be tested not only on new images, but on entirely unseen instances of the material.

Each sample contributes 81 images to a class, generated under a combination of three illumination conditions, three viewing angles, and three different scales. As a result, each class has a total of 324 images (81 images \times 4 samples), leading to a complete dataset of 3564 images across all 11 classes. The original image sizes vary around 200×200 pixels, but for training consistency, they are uniformly resized to 128×128 pixels. The images are color (RGB), and preprocessing steps such as CLAHE (Contrast Limited Adaptive Histogram Equalization) are applied to improve local texture contrast without distorting the structural integrity of the material surfaces.

15 To create a reliable and unbiased model evaluation setup, the dataset is split in a way that mirrors a domain generalization problem. The samples *sample_a*, *sample_b*, and *sample_c* are used collectively to form the training and validation set, while *sample_d* is held out exclusively for testing. This ensures that the test data comes from a completely unseen sample, allowing a realistic assessment of the model's ability to generalize to new material instances captured under unseen conditions.

4 Within the training set, to maintain a balanced and fair evaluation across all classes, we employ **stratified K-fold cross-validation** during the training-validation split. This technique ensures that each fold retains the same proportion of class distributions as in the original dataset. Since each of the three training samples contributes 81 images per class, we have 324 images per class for training and validation, resulting in 3564 images total (324×11 classes). Stratified K-fold splitting helps to distribute these images such that both the training and validation subsets contain a representative and balanced mix of all classes. This is crucial, especially in texture classification, where some classes may share visual similarities, and imbalanced data could lead to biased learning.

A typical 80-20 train-validation split under this scheme would yield approximately 260 images per class for training and 49 for validation. This corresponds to around 2860 training images and 712 validation images across all classes. The testing set, derived solely from *sample_d*, contains 81 images per class, summing up to a total of 972 test images.

16 By adopting this sample-aware and stratified data preparation strategy, we ensure that our models are trained on a rich, balanced distribution of textures while being rigorously evaluated on genuinely unseen material instances. This not only helps in preventing overfitting but also enhances the model's ability to generalize across scale, lighting, and viewpoint variations — which are critical for real-world texture recognition tasks.

3.4.2 Implementation

The implementation and hyperparameter tuning of the three models—pretrained ResNet18, pretrained ResNet50, and ResNet50 with histogram layers—played a crucial role in achieving effective texture classification on the KTH-TIPS2a dataset. Each model presented its own architectural capacity and sensitivity to tuning, and their performance depended heavily on appropriate regularization, learning dynamics, and alignment with the nature of texture variation in the dataset. **Pretrained ResNet18 model** The initial focus of the pipeline is on preparing the KTH-TIPS2a dataset. Since texture datasets often involve variations in scale, lighting, and viewpoint, special care is taken to enhance local features while preserving texture integrity. To this end, a Contrast Limited Adaptive Histogram Equalization (CLAHE) technique is applied to each image. This transformation enhances local contrast in a controlled way and ensures that small texture variations are more distinguishable without blowing out global intensity variations.

Next, two different transformation pipelines are set up—one for training and another for validation/testing. The training transform is aggressive and diverse, designed to increase the generalization power of the model. It includes random rotations, flips (both horizontal and vertical), and color jitter to simulate different lighting conditions. Gaussian blur and Gaussian noise are added selectively to model variability in real-world textures. All training images are finally normalized to have a mean and standard deviation of 0.5 in each channel, and resized to a fixed input size. In contrast, the validation and test images are transformed more conservatively using only CLAHE, resizing, and normalization, ensuring consistent evaluation without introducing noise or distortion.

Once the dataset is defined, it is split into a training and a validation set using an 80:20 ratio. DataLoaders are created for all subsets to handle batching, shuffling (in training), and parallel loading.

26 It served as the baseline due to its lightweight architecture and proven effectiveness on general vision tasks. Being a relatively shallow network with fewer parameters, it was suitable for initial experimentation on a limited dataset like KTH-TIPS2a. The model was initialized with ImageNet-pretrained weights and fine-tuned on the texture data after replacing the final classification layer to output 11 class scores. In training, learning rate was a key hyperparameter—initial experiments with a learning rate of 1e-3 led to unstable loss and poor generalization, so it was reduced to 5e-4, which provided a stable and smooth convergence. Optimizers like Adam and SGD with momentum were explored, and while Adam led to faster initial learning, SGD gave slightly more stable performance across epochs when early stopping was applied. A batch size of 16 was used, balancing training stability and GPU memory constraints. Overall, while ResNet18 captured low-to mid-level texture features well, its limited depth posed difficulties in resolving complex textures under extreme illumination and scale variation.

35 The pretrained ResNet50 model marked a significant improvement in performance, owing to its deeper architecture and increased capacity to model hierarchical features. The model was also initialized with ImageNet weights and the final layer modified for 11-class classification. Since ResNet50 contains many more parameters and deeper convolutional layers, the learning rate was carefully lowered to 1e-4 to avoid disrupting the pretrained filters during fine-tuning. The first few layers of the network were initially frozen for stability, and then gradually unfrozen after about ten epochs to allow full backpropagation through the network. AdamW was preferred over regular Adam due to its decoupled weight decay, which improved regularization in deep networks. The AdamW optimizer is an advanced optimization technique used in training deep learning models, especially effective in complex neural architectures such as transformer-based networks and convolutional neural networks. It builds upon the popular Adam optimizer, which combines the benefits of both momentum-based and adaptive learning rate strategies. While Adam estimates the first and second moments of the gradients to adjust learning rates per parameter, it also includes a weight decay mechanism, typically implemented as L2 regularization. However, in standard Adam, this weight decay is directly added to the gradient, thereby entangling the regularization effect with the adaptive learning rate logic. The training required smaller batch sizes, often 32, and longer training schedules up to 100 epochs due to the complexity of the network. ResNet50's architecture, with its residual connections and wider convolutional filters in the later stages, proved particularly useful in learning texture patterns across varied lighting and viewpoints, which were not well captured by the baseline model.

Weight decay is a regularization technique used in training machine learning models—particularly neural networks—to prevent overfitting by discouraging overly large weights. At its core, weight decay operates by adding a penalty to the model’s loss function that increases as the magnitude of the model’s weights increases. The goal is to keep the learned weights relatively small, which helps the model generalize better to unseen data.

The idea behind this stems from the observation that models with large parameter values often tend to fit the training data too closely, capturing noise and irregular patterns rather than the underlying structure. This phenomenon, known as overfitting, leads to poor performance on validation or test sets. By encouraging smaller weights, weight decay biases the model toward simpler solutions that are less likely to memorize the training set and more likely to generalize well. To further enhance the model’s ability to differentiate textures under challenging conditions, histogram layers were introduced into the ResNet50 architecture. The histogram-based model was designed to capture global statistical distributions of intermediate features, which are often crucial in distinguishing fine textures with subtle local variations or ambiguous patterns. These histogram layers were added after the third residual block of ResNet50, where feature maps are sufficiently abstract but still retain spatial structure. Instead of simply relying on deep convolutional features, the histogram layer aggregated feature activations into learned bins, capturing frequency-like statistics that are useful in identifying repetitive or stochastic patterns in textures such as those found in cork, aluminum foil, and brown bread.

This architectural enhancement, however, required careful hyperparameter tuning. A much smaller learning rate of $5e-5$ was necessary to prevent over-updating the histogram bins, which were sensitive to gradient changes. The training also demanded more epochs, typically 70 or more, as histogram layers converged more slowly than standard convolutional layers. The optimizer was kept as AdamW to maintain weight regularization. Because histogram layers increase the model’s expressiveness, dropout regularization and data augmentation were crucial to avoid overfitting. These layers helped the model better capture the statistical diversity of the texture classes and proved particularly effective in challenging classes like aluminum foil, cracker, and white bread where texture variance is high across scales and lighting.

Table 3.1: Hyperparameter Settings for Each Model

Hyperparameter	ResNet18	ResNet50	ResNet50 + Histogram Layers
Pretraining	ImageNet	ImageNet	ImageNet
Learning Rate	5×10^{-4}	1×10^{-4}	5×10^{-5}
Optimizer	Adam / SGD	AdamW	AdamW
Batch Size	32	32	32
Epochs	48	60	65
Initial Freezing	Unfrozen	Unfrozen	Unfrozen
Dropout	0.2	0.2	0.2
Histogram Bins	–	–	4
Data Augmentation	Flip, Rotate	Flip, Rotate, ColorJitter	Extensive (Flip, Rotate, Zoom, Brightness)
Regularization	Weight Decay ($1e-4$)	Weight Decay ($1e-4$)	Weight Decay ($1e-4$)

The implementation and tuning of each model were tailored to the dataset's specific demands. ResNet18 provided a solid baseline but lacked the depth to generalize across wide variability. ResNet50 offered better depth and residual learning, enabling improved performance on more complex texture features. Finally, the histogram-enhanced ResNet50 capitalized on both spatial and statistical cues, significantly enhancing the model's ability to generalize across the dataset's intricate texture classes when paired with carefully tuned learning parameters.

3.5 Results of the Models

3.5.1 Pretrained Resnet18 Results

The finetuned **ResNet18** model trained on the **KTH-TIPS2a texture classification dataset** demonstrates strong performance on the training and validation sets but reveals important patterns of overfitting and generalization challenges when evaluated on the test set.

From the training statistics, the model achieves an **exceptionally low training loss of 0.0806** and a **training accuracy of 98.81%**, suggesting it has effectively learned patterns from the training data. The **validation loss is extremely low (0.0003)** with a **perfect validation accuracy of 100%**, which might seem ideal on the surface but actually raises concerns of **overfitting**—especially when juxtaposed with the **test loss of 0.9964** and a **test accuracy of 81.99%**. This large gap between validation and test performance clearly indicates that the model may be **overly tuned to the training/validation distribution**, likely memorizing patterns that do not generalize well to unseen data.

Delving into the **classification report**, the overall test **accuracy of 82%** aligns with this conclusion. Although the model performs well on many classes, it shows clear disparities across different textures, especially where visual features are subtle or overlapping. The **macro average** metrics—**precision (0.83)**, **recall (0.82)**, and **f1-score (0.81)**—suggest balanced performance across classes without being skewed by class size. Meanwhile, the **weighted averages**—**precision (0.85)** and **f1-score (0.83)**—indicate that the model performs especially well on more frequently occurring or easier-to-recognize classes. This balance confirms that while the model captures broad patterns effectively, there remains room for improvement in generalizing to more difficult texture classes.

Examining individual class performance reveals the source of this imbalance. The model shows **outstanding classification** for some textures—**white_bread** (precision: 0.98, recall: 0.97, f1-score: 0.98), **cork** (f1-score: 0.92), **wood** (0.92), and **lettuce_leaf** (0.93)—indicating that these textures likely have distinctive and consistent patterns that the convolutional filters are able to learn well.

In contrast, textures like **cotton** (f1-score: 0.53), **cracker** (0.70), and **wool** (0.73) exhibit significantly lower scores. For instance, **cotton** has **very low precision (0.45)**, meaning that many images from other classes are being falsely labeled as cotton. **Cracker** shows the opposite issue, with high recall (0.94) but low precision (0.56), indicating over-prediction of this class. These behaviors point to **poor feature separation** or **class confusion**, likely due to similar low-contrast or repetitive patterns shared with other classes like wool, linen, or corduroy.

Some mid-performing classes like **corduroy** (f1-score: 0.78) and **linen** (0.79) further support the idea that **inter-class similarity** and **intra-class variability** remain persistent challenges for this model.

Classification Report:				
	precision	recall	f1-score	support
aluminium_foil	1.00	0.76	0.87	72
brown_bread	0.95	0.70	0.81	108
corduroy	0.84	0.72	0.78	108
cork	0.92	0.92	0.92	108
cotton	0.45	0.62	0.53	72
cracker	0.56	0.94	0.70	72
lettuce_leaf	0.88	0.98	0.93	108
linen	0.77	0.81	0.79	72
white_bread	0.98	0.97	0.98	108
wood	0.99	0.86	0.92	108
wool	0.80	0.68	0.73	108
accuracy			0.82	1044
macro avg	0.83	0.82	0.81	1044
weighted avg	0.85	0.82	0.83	1044

Figure 3.8: Classification report of Resnet18

The presence of **high training and validation scores but noticeably lower test accuracy** supports the hypothesis that the model has not fully generalized. Overfitting may stem from a lack of texture diversity in the training data or insufficient regularization. This makes clear that strong in-sample performance does not necessarily translate to real-world generalization, especially in texture classification tasks where **illumination, scale, and orientation** can greatly affect appearance.

3.5.2 Pretrained Resnet50 Results

The finetuned ResNet50 model demonstrates a clear improvement over the previous ResNet18 implementation, achieving 86.21% test accuracy with a test loss of 0.7881. While this shows the deeper architecture's better generalization capabilities, the significant gap between perfect validation accuracy (99.86%) and test performance confirms that overfitting remains an issue - though less severe than in ResNet18's case where we observed an 18% discrepancy compared to the current 13% gap.

The classification report reveals several important patterns in the model's behavior. It excels with high-contrast, structurally distinct textures like lettuce_leaf (perfect f1-score of 1.00) and white_bread (0.96), where convolutional filters can easily identify unique patterns. However, fibrous textures such as cotton (f1-score: 0.54) and linen (0.79) continue to pose significant challenges, with their interwoven patterns leading to frequent misclassifications that reveal the limits of hierarchical feature learning.

Particularly interesting is cork's performance trajectory: it achieves perfect recall (1.00) but suffers from precision drops (0.72), indicating the model now detects all cork samples but often mislabels other crinkled materials as cork. Brown_bread shows the opposite tendency - high precision (0.99) but modest recall (0.65) - suggesting the model is conservative in assigning this label, potentially confusing darker bread samples with similarly colored crackers.

Comparative Analysis with ResNet18

When compared directly with the ResNet18 implementation, the ResNet50 model shows measurable improvements that validate the benefits of increased network depth. The 4.22% boost in test accuracy (from 81.99% to 86.21%) and reduction in test loss (from 0.9964 to 0.7881) confirm that additional layers contribute to better generalization. The

Classification Report:				
	precision	recall	f1-score	support
aluminium_foil	1.00	0.83	0.91	72
brown_bread	0.99	0.65	0.78	108
corduroy	0.78	0.82	0.80	108
cork	0.72	1.00	0.84	108
cotton	0.63	0.47	0.54	72
cracker	0.80	0.97	0.88	72
lettuce_leaf	0.99	1.00	1.00	108
linen	0.88	0.72	0.79	72
white_bread	0.95	0.97	0.96	108
wood	0.91	0.96	0.94	108
wool	0.86	0.93	0.89	108
accuracy			0.86	1044
macro avg	0.87	0.85	0.85	1044
weighted avg	0.87	0.86	0.86	1044

Figure 3.9: Classification report of Resnet50

narrowing train-test accuracy gap (from 18% to 13%) suggests ResNet50 is more resistant to overfitting, though not immune to it.

Class-specific comparisons reveal where increased depth helps most. Cork shows an f1-score improvement from 0.84 to 0.92, cracker from 0.70 to 0.88, and wool from 0.73 to 0.89 - suggesting ResNet50's additional convolutional blocks better capture multi-scale patterns in these materials. However, some challenges remain stubbornly unchanged - cotton's f1-score only improves marginally from 0.53 to 0.54, indicating neither architecture adequately resolves fine discrimination of similar woven fabrics.

The evolution of precision-recall tradeoffs between architectures is particularly noteworthy. In ResNet18, cracker showed high recall (0.94) but low precision (0.56), indicating overprediction. ResNet50 maintains the high recall (0.97) while significantly improving precision (0.80), demonstrating better calibration. Meanwhile, linen's performance remains nearly identical between architectures (0.79 f1-score in both), suggesting this particular texture confusion may represent a fundamental limitation of convolutional approaches regardless of depth.

The comparative results suggest several important implications for model development. ResNet50's across-the-board improvements validate the value of increased depth for texture analysis, particularly for materials with multi-scale patterns. However, the persistent struggles with certain texture categories indicate that simply adding more convolutional layers may not solve all material recognition challenges.

The consistent performance gaps between training and test metrics in both architectures highlight the need for better regularization strategies. While ResNet50's deeper structure provides some inherent regularization, additional techniques like aggressive data augmentation or hybrid approaches combining CNNs with handcrafted texture features may be necessary to bridge the remaining generalization gap.

The model's varied performance across texture categories is particularly revealing. Strong results for structured, inorganic materials (like aluminium foil) versus weaker performance on organic, variable textures (like cotton or linen) suggests that future architectures might benefit from domain-specific adaptations. For fibrous materials, incorporating directional filters or spectral analysis components could potentially capture the periodic patterns that current convolutional operations miss.

3.5.3 Pretrained Resnet50 with Histogram Layers Results

The ResNet50 architecture augmented with histogram layers demonstrates exceptional performance on the KTH-TIPS2a texture classification task, achieving a test accuracy of 99.78% with a remarkably low test loss of 0.0145. This represents a substantial improvement over the standard ResNet50 model's 86.21% accuracy and 0.7881 test loss observed in previous experiments. The model achieves near-perfect classification across nearly all texture categories, as evidenced by the classification report showing perfect 1.00 precision, recall, and f1-scores for all classes except wool, which achieves a slightly lower but still outstanding 0.99 f1-score.

The training metrics reveal equally impressive results, with 98.27% training accuracy (loss: 0.1544) and 98.51% validation accuracy (loss: 0.0702). Unlike previous implementations that showed significant gaps between training and test performance, this model maintains exceptional consistency across all phases of evaluation. The minimal discrepancy between training (98.27%) and test (99.78%) accuracy indicates the histogram layers contribute to fundamentally more robust feature learning compared to the standard ResNet50 architecture.

Where previous models struggled with characteristic failure patterns - particularly in discriminating between similar fibrous textures like cotton, linen, and wool - this implementation demonstrates flawless performance. The perfect 1.00 f1-scores for previously problematic classes like cotton (compared to 0.54 in standard ResNet50) and corduroy (compared to 0.80) suggest the histogram features provide precisely the complementary information needed to resolve ambiguities that challenged pure convolutional approaches. Even wool's slight imperfection (0.99 f1-score due to 0.98 recall) represents a dramatic improvement over previous results (0.89 in standard ResNet50, 0.73 in ResNet18).

The model's performance represents several significant advances over previous architectures. First, it eliminates the precision-recall tradeoffs that were evident in earlier implementations - where improvements in one metric often came at the expense of the other. Second, it demonstrates unprecedented consistency across all material categories, solving the previous "texture-type bias" where models performed well on structured inorganic materials but struggled with organic patterns. Third, it achieves these results while maintaining balanced training dynamics, as shown by the harmonious alignment of train/val/test metrics that suggests genuine learning rather than memorization.

Classification Report:				
	precision	recall	f1-score	support
aluminium_foil	1.00	1.00	1.00	79
brown_bread	1.00	1.00	1.00	86
corduroy	1.00	1.00	1.00	86
cork	1.00	1.00	1.00	87
cotton	0.98	1.00	0.99	79
cracker	1.00	1.00	1.00	79
lettuce_leaf	1.00	1.00	1.00	86
linen	1.00	1.00	1.00	79
white_bread	1.00	1.00	1.00	87
wood	1.00	1.00	1.00	87
wool	1.00	0.98	0.99	87
accuracy			1.00	922
macro avg	1.00	1.00	1.00	922
weighted avg	1.00	1.00	1.00	922

Figure 3.10: Classification report of pretrained Resnet50 with histogram layers

The remaining 0.22% error margin, primarily from wool’s occasional misclassifications, indicates the model’s current limitations. Unlike previous architectures that failed categorically on whole texture groups, this model’s shortcomings appear more nuanced, likely involving rare cases of extreme lighting variations or material deformations that affect both spatial and statistical texture signatures. This progression from gross failures to subtle edge cases marks a significant advancement in the technology’s capabilities.

These results fundamentally change our understanding of effective texture analysis architectures. Where increasing network depth alone (from ResNet18 to ResNet50) brought only incremental improvements, the integration of histogram features with convolutional operations has produced transformative gains. The implications extend beyond this specific dataset, suggesting that combining structural and statistical representations may offer similar breakthroughs in other fine-grained material recognition tasks where traditional CNNs have previously plateaued.

Model Variant	Training Loss	Validation Loss
Finetuned ResNet18	0.0806	0.0003
Finetuned ResNet50	0.0555	0.0132
ResNet50 + Hist Layers	0.1544	0.0702

Table 3.2: Training and Validation Loss Comparison

Model Variant	Validation Accuracy (%)	Test Accuracy (%)
Finetuned ResNet18	100.00	81.99
Finetuned ResNet50	99.86	86.21
ResNet50 + Hist Layers	98.51	99.78

Table 3.3: Validation and Test Accuracy Comparison

Chapter 4

Conclusion

This study addressed texture classification on the challenging KTH-TIPS2a dataset using deep learning. Starting with a pretrained ResNet18 baseline, we observed limitations in generalization due to high intra-class variability and limited data. Improvements were achieved with a deeper ResNet50 and further enhanced using Histogram Layers, which helped capture fine-grained texture patterns. Effective preprocessing and augmentation were key to handling variations in scale, lighting, and viewpoint. Overall, combining deep features with texture-specific encodings proved effective for robust texture recognition. Further enhancement was achieved by integrating Histogram Layers into the ResNet50 architecture. This allowed the model to explicitly capture local texture distributions, particularly beneficial for high-frequency or low-contrast materials like aluminum foil and cotton. Careful image preprocessing, including CLAHE and strong data augmentations, played a crucial role in mitigating illumination and viewpoint biases. Finally, through detailed analysis of misclassified points and architectural improvements, we demonstrated that combining deep hierarchical features with statistical texture encodings offers a robust solution for fine-grained texture recognition, especially under real-world imaging conditions.

Bibliography

- [1] K. He, X. Zhang, S. Ren, and J. Sun, "Deep Residual Learning for Image Recognition," *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 770–778, 2016. arXiv:1512.03385.
- [2] K. He, X. Zhang, S. Ren, and J. Sun, "Identity Mappings in Deep Residual Networks," *European Conference on Computer Vision (ECCV)*, pp. 630–645, 2016. arXiv:1603.05027.
- [3] S. Xie, R. Girshick, P. Dollár, Z. Tu, and K. He, "Aggregated Residual Transformations for Deep Neural Networks," *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 1492–1500, 2017. arXiv:1611.05431.
- [4] Y. Bao, L. Zhang, and X. Wang, "A Survey on Deep Residual Networks," *Applied Sciences*, vol. 10, no. 5, p. 1533, 2020. doi:10.3390/app10051533.
- [5] J. Peeples, W. Xu, and A. Zare, "Histogram Layers for Texture Analysis," *IEEE Transactions on Artificial Intelligence*, vol. 3, no. 4, pp. 541–552, 2022. DOI: 10.1109/TAI.2021.3135804. :contentReference[oaicite:1]index=1
- [6] Z. Wang, H. Li, W. Ouyang, and X. Wang, "Learnable Histogram: Statistical Context Features for Deep Neural Networks," in *European Conference on Computer Vision (ECCV)*, 2018. :contentReference[oaicite:2]index=2
- [7] J. Peeples, A. Zare, J. Dale, and J. Keller, "Histogram Layers for Synthetic Aperture Sonar Imagery," arXiv preprint arXiv:2209.03878, 2022. :contentReference[oaicite:3]index=3
- [8] J. Peeples, S. Al Kharsa, L. Saleh, and A. Zare, "Histogram Layers for Neural Engineered Features," arXiv preprint arXiv:2403.17176, 2024. :contentReference[oaicite:4]index=4
- [9] J. Yosinski, J. Clune, Y. Bengio, and H. Lipson, "How transferable are features in deep neural networks?" *Advances in Neural Information Processing Systems (NeurIPS)*, vol. 27, 2014.
- [10] M. Huh, P. Agrawal, and A. A. Efros, "What makes ImageNet good for transfer learning?" *arXiv preprint arXiv:1608.08614*, 2016.
- [11] C. Tan, F. Sun, T. Kong, W. Zhang, C. Yang, and C. Liu, "A Survey on Deep Transfer Learning," *International Conference on Artificial Neural Networks (ICANN)*, pp. 270–279, 2018. arXiv:1808.01974.

- 10 [12] N. Tajbakhsh et al., “Convolutional Neural Networks for Medical Image Analysis: Full Training or Fine Tuning?” *IEEE Transactions on Medical Imaging*, vol. 35, no. 5, pp. 1299–1312, 2016.
- 9 [13] R. M. Haralick, K. Shanmugam, and I. Dinstein, “Textural Features for Image Classification,” *IEEE Transactions on Systems, Man, and Cybernetics*, vol. SMC-3, no. 6, pp. 610–621, 1973. :contentReference[oaicite:1]index=1
- 18 [14] P. Porebski and M. Vandenbroucke, “Haralick Feature Extraction from LBP Images for Color Texture Classification,” *Journal of Visual Communication and Image Representation*, 2018. :contentReference[oaicite:6]index=6