

THIR: Topological Histopathological Image Retrieval

Anonymous Full Paper
Submission 8

Abstract

According to the World Health Organization, breast cancer claimed the lives of approximately 685,000 women in 2020. Early diagnosis and accurate clinical decision-making are critical in reducing this global burden. In this study, we propose THIR, a novel Content-Based Medical Image Retrieval (CBMIR) framework that leverages topological data analysis—specifically, Betti numbers derived from persistent homology—to characterize and retrieve histopathological images based on their intrinsic structural patterns. Unlike conventional deep learning approaches that rely on extensive training, annotated datasets, and powerful GPU resources, THIR operates entirely without supervision. It extracts topological fingerprints directly from RGB histopathological images using cubical persistence, encoding the evolution of loops as compact, interpretable feature vectors. The similarity retrieval is then performed by computing the distances between these topological descriptors, efficiently returning the top- K most relevant matches.

Extensive experiments on the BreakHis dataset demonstrate that THIR outperforms state-of-the-art supervised and unsupervised methods. It processes the entire dataset in under 20 minutes on a standard CPU, offering a fast, scalable, and training-free solution for clinical image retrieval.

1 Introduction

Cancer is a leading cause of death worldwide, with nearly 10 million deaths in 2020, or nearly one in six deaths [1]. Breast cancer accounts for 25% of all cancers in women worldwide, and about 685,000 women lost their lives due to breast cancer in 2020 [2]. Histopathology is the gold standard for cancer diagnosis [3], which involves extracting tissue specimens from suspicious areas to prepare a glass slide for a microscopic examination [4]. However, this examination might have some human errors or intraobserver variability. Singh, et al. [5] made an extensive review of the errors in cancer diagnosis. Accurate cancer diagnosis and grading rely on many factors, including the knowledge, experience, and skills of pathologists [6], which can increase the rate of human error in diagnosis. Diagnosis is a high-risk area of errors, including missed, inappropriately delayed, or wrong diagnoses [6]. Digital pathology, by

employing Deep Learning (DL) and Machine Learning (ML) techniques, has a significant impact on decreasing human errors by providing a second opinion for pathologists [7]. An image retrieval tool that finds cases with similar morphological features can help diagnose rare diseases and unusual conditions that may not have enough cases available to develop accurate supervised classification models [8].

While Content-Based Image Retrieval (CBIR) has been under investigation for decades [9], only with the emergence of digital pathology and DL, the studies have begun to focus on image search and analysis in histopathology [10]. Content-Based Medical Image Retrieval (CBMIR) offers a new approach to computational pathology [11]. CBMIR provides the top K with patches similar to the query from the previously diagnosed and treated cases. This can assist pathologists in tackling the above-mentioned errors. The main base of CBMIR is similarity measurement, which considers features including texture, shape, intensity, etc., and compares them with the previous cases [12]. This retrieval helps pathologists receive not only the labels, but also patches similar to their query. This can increase the explainability of the methods since pathologists can analyze the texture and compare it with similar cases based on their expertise.

Visual examination of the patterns of the tissue in a monitor is a task usually relegated to pathologists or computational biologists. CBMIR can assist pathologists in analyzing and managing a large volume of images to enhance diagnosis, collaboration, education, research, and the decision-making processes [13]. By searching and retrieving visually similar images with their labels, pathologists have more information to detect the abnormalities. CBMIR not only increases the accuracy of cancer diagnosis but also speeds up the process of consulting peers. It can support remote consultation and collaboration between pathologists from all over the world [14]. In addition, CBMIR is a crucial tool in research for exploring large histopathological image archives to discover and identify new patterns, trends, graphs, and correlations between cancer grades [13]. For example, it can be a teaching tool in histopathological education, allowing trainers to study and compare various histopathological patterns.

CBMIR consists of ranking images concerning a query image based on visual similarities with a typical workflow, as illustrated in Figure 1. Following [15], it has two phases: offline and online.

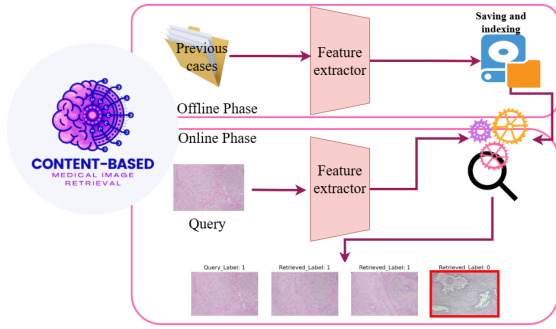


Figure 1. shows the main workflow of a CBMIR. It contains two main phases, offline and online. The same FE applies in both phases to extract features. Then, the Euclidean distance as a distance measurement function is applied to find the top-3 similar patches. On top of the query image and the retrieved images, their labels are mentioned.

999 The offline phase includes extracting features of the
1000 previous cases and indexing them. In the online
1001 phase, a query image, which is an unseen image
1002 for the Feature Extractor (FE), is fed to the same
1003 FE as in the offline phase to extract its features.
1004 To identify those most relevant images, a similarity
1005 function is applied between the extracted features
1006 of the query and the previous cases. Then, the top- K
1007 most relevant images from the previous cases are
1008 retrieved. This is similar to the traditional workflow
1009 in hospitals using Atlas books [16].

110 DL-based methods, particularly Convolutional
111 Neural Networks (CNNs), have shown great suc-
112 cess in CBMIR tasks by automatically learning hi-
113 erarchical feature representations. However, these
114 models often require large amounts of labeled data,
115 are prone to overfitting, suffer from limited GPU
116 resources, and lack interpretability. In this paper,
117 THIR focuses on the topological information of im-
118 ages and texture features extracted from Topological
119 Data Analysis (TDA). TDA provides hand-crafted,
120 mathematically interpretable features, such as Betti
121 values, that describe the global structure of images.
122 While DL focuses on local patterns and appearance,
123 TDA emphasizes global connectivity and shape, of-
124 fering complementary information [17].

125 To the best of the authors’ knowledge, this is
126 the first study to apply TDA to CBMIR in digital
127 pathology. The main contributions of this paper are
128 as follows:

- 129 • **THIR** is a fully unsupervised method that elim-
130 inates the need for labeled datasets, addressing
131 one of the key challenges in medical image anal-
132 ysis.
- 133 • Through cubical persistence, we capture unique
134 topological signatures by tracking how topologi-
135 cal structures develop across the color channels
136 of the images.
- 137 • **THIR** is computationally efficient; the entire

feature extraction process for the training set 138
using THIR takes approximately 20 minutes, 139
resulting in a compact and informative topolog- 140
ical feature vector. 141

- Unlike ML-/DL-based methods, **THIR** does 142
not require hyperparameter tuning, GPU re- 143
sources, and a training phase, making it simple, 144
scalable, and accessible. 145
- The accuracy of **THIR** surpasses that of re- 146
cently published supervised and unsupervised 147
methods in CBMIR tasks. 148

2 Related work 149

The performance of CBMIR mainly depends on the 150
choice and performance of the FE method. A high- 151
quality FE algorithm can improve the precision of 152
the search engine [18]. There are some common 153
descriptors, such as Scale-invariant feature trans- 154
form (SIFT) [19], Local Binary Patterns (LBP) [20], 155
Histogram of Oriented Gradient (HOG) [21], edge 156
histogram descriptor [22], and Gabor filter [23] for 157
texture analysis, which explore the local features 158
of the images. Local features refer to the general 159
pattern of images, such as a point, edge, or small 160
image patch. 161

Artificial intelligence (AI) enabled CBMIR for an 162
effective diagnosis in [24]. Then, as CNNs show their 163
power and high effectiveness in extracting features, 164
they have become increasingly used for CBMIR. DL- 165
based models yield high-performance search engines 166
by extracting features of images for tasks related to 167
CBMIR. Many recent studies [25–28] were dedicated 168
to exploring the performance of different DL-based 169
methods in CBMIR. 170

Among the various types of DL-based methods, 171
Auto Encoders (AEs), GANs, and Siamese net- 172
works [29] have a special place as FE in the CBMIR 173
task. In [14], the author reported 9.33 and 6.59 hours 174
of training time for training a Convolutional Auto 175
Encoder (CAE) and Federated Learning (FL) CAE, 176
respectively. [14] claims that FedCBMIR is faster 177
compared to traditional CBMIR using the same 178
CAE structure and the same GPU type (NVIDIA 179
GeForce RTX 3090). FedCBMIR provides 98% accu- 180
racy, and the UCBMIR in [30] yields 93% accuracy 181
at the top-5 on the BreakHis data set. The Siamese 182
network in [31] obtains 94% an F1-score at the top- 183
5 retrievals for breast cancer. Authors in [32] in 184
Google AI Healthcare proposed an automatic high- 185
level feature extraction on prostate cancer. The 186
obtained results were reported at the top-5 simi- 187
lar patches with an accuracy of 73%. Yottixel [33] 188
uses the DenseNet structure, which is trained on 189
the ImageNet data set for extracting patches with- 190
out being trained specifically for the CBMIR task. 191

192 RetCCL [34] proposes a method based on cluster-
193 ing feature vectors of the patches. In this work,
194 a ResNet50 was trained using contrastive learning.
195 In [35], a Graph Neural Network (GNN) encodes
196 Region of Interest (ROI) graphs into representations
197 using a contrastive loss function in a self-supervised
198 manner. The study in [36] proposes size-scalable
199 CBMIR from databases that consist of whole-slide
200 images (WSIs). This method has addressed scal-
201 able retrieval frameworks tailored to WSIs, focus-
202 ing on efficient indexing and patch-level compar-
203 isons to manage the immense size and complexity
204 of histopathological data.

205 The authors in [37] provide an overview of the
206 TDA methods in biomedicine. After reviewing the
207 recently published literature, this study aims to
208 explore the potential of TDA with CBMIR applica-
209 tions. To the best of our knowledge, this combina-
210 tion has never been explored for CBMIR tasks in
211 breast cancer.

212 3 Material and Methodology

213 Our methodology consists of two steps. First, we
214 extract the topological feature vectors from the data
215 set and the query. Then, we apply a similarity
216 measurement function to these vectors to find the
217 top- K similar patches to the query from the data
218 set. This provides a search engine based on the
219 images’ topological information, resulting in a fast
220 and interpretable model called THIR. The imple-
221 mentation in this study focuses on breast cancer.
222 Figure 2 and Figure 3 illustrate the whole workflow
223 of the proposed methodology.

224 3.1 Data set

225 BreakHis data set [38] was created in the PD lab-
226 oratory in Prana, Brazil, and consists of 7909 mi-
227 croscopic images of breast cancer. This collection
228 contains four different magnifications ($40\times$, $100\times$,
229 $200\times$, and $400\times$)¹ as shown in Table 1. In this bi-
230 nary data set, tissues were stained with Hematoxylin
231 and Eosin (H&E), which is the most common color
232 in histopathological images [39]. Following the pre-
233 vious studies [14, 30, 31] to be able to have compa-
234 rable results, we resized the images into $240\times 240\times 3$.
235 Since the images are from the cancerous tissue, they
236 are not affected by image transforming, inverting,
237 zooming in, or rotation by 90 degrees [40].

238 3.2 Topological Data Analysis (TDA)

239 Topology studies properties of spaces that are in-
240 variant under any continuous deformation. Over the
241 past two decades, TDA has proven highly effective in
242 identifying topological structures within images [41].

¹<https://www.kaggle.com/datasets/ambarish/breakhis>

Table 1. The distribution of BreakHis data set.

Magnification	Benign	Malignant	Total
40×	625	1370	1995
100×	644	1437	2081
200×	623	1390	2013
400×	588	1232	1820
Total	2480	5429	7909

243 In the context of histopathological image analysis,
244 TDA has demonstrated strong potential in cancer
245 detection [42]. TDA extracts meaningful patterns by
246 analyzing the homological features of images. These
247 features can quantify the complex topological shapes
248 and geometric structures in the data. The advantage
249 of TDA is that it can effectively process complex
250 and high-dimensional data, capture the global topo-
251 logical structure of the data, and provide a deep
252 understanding of the shape of the data [43]. This
253 means that TDA offers a mathematically robust and
254 interpretable way to capture topological features,
255 especially in medical images where texture, shape,
256 and structure matter.

257 In this study, we utilize TDA to extract topologi-
258 cal features from medical images, specifically fo-
259 cusing on Persistent Homology (PH), one of TDA’s
260 most prominent tools. PH quantifies how topologi-
261 cal features—such as connected components, loops,
262 and voids—evolve across different scales, represented
263 by a filtration. A filtration is a sequence of nested
264 spaces generated by progressively thresholding the
265 data. Within this framework, a feature *is born* at
266 the threshold where it first appears and *dies* at the
267 threshold where it is merged or disappears. Long-
268 lived features (those with large persistence) are con-
269 sidered topologically meaningful, while short-lived
270 ones are often attributed to noise [44]. A compre-
271 hensive overview of TDA and PH is provided in [45].

272 There are two widely used algorithms for compu-
273 ting homology: simplicial and cubical homology.
274 While simplicial complexes offer greater generality,
275 cubical complexes are more computationally efficient
276 and well-aligned with image data, which naturally
277 reside on regular pixel grids [46, 47]. In the cubical
278 setting, images are modeled as 2D grids of square
279 cells (pixels), and a topological structure is built by
280 thresholding pixel intensities. The resulting binary
281 images are interpreted using 0D (points), 1D (edges),
282 and 2D (squares) elements [48].

283 As the threshold increases, new topological fea-
284 tures appear and disappear, which can be quantified
285 using Betti numbers: β_0 (number of connected com-
286 ponents), β_1 (number of loops), and β_2 (number of
287 voids). These values provide compact, interpretable
288 descriptors of image structure and texture. In medi-
289 cal imaging, they serve as powerful shape-based
290 features that are invariant to rotation and inten-
291 sity changes and are equivalent to spatial scaling,
292 offering a complementary perspective to pixel-based

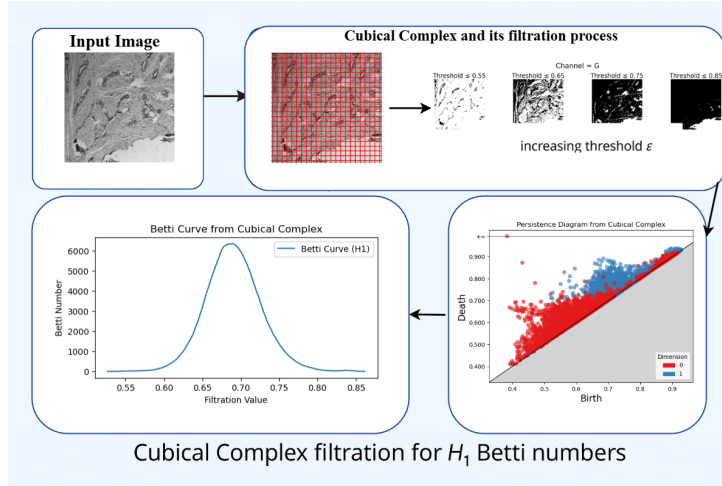


Figure 2. THIR model. We first generate persistence diagrams for any input images, utilizing the cubical complex on each channel of the images. Next, we derive our topological feature vectors, represented as the Betti curves. The values of this curve are then input into the CBMIR workflow to produce the results of the search engine.

293 or deep learning methods. This interpretability is
294 particularly valuable in clinical contexts, where ex-
295 plainability is essential.

296 Figure 2 illustrates the cubical complex process.
297 The input image is first overlaid with a grid, and
298 each unique intensity value is treated as a threshold
299 value (th). For 8-bit grayscale images, the filtration
300 spans $K_r r \in \mathcal{T} = K_0, \dots, K_{255}$. As the th value
301 increases, connected components and loops emerge
302 and then merge or vanish. The persistence diagram
303 in the figure summarizes this process: β_0 features
304 are shown in red and β_1 in blue. The diagonal line
305 represents $birth = death$, and points far from the
306 diagonal indicate more persistent, and thus more
307 meaningful features [49].

308 In this study, we focus on β_1 . So, the Betti values
309 are the number of loops in the image. As can be
310 seen in Figure 2, the Betti curve for this data set
311 resembles a bell shape, peaking where most loops
312 are present. Initially, at low thresholds, few fea-
313 tures appear due to limited pixel activation. As
314 the threshold increases, more features emerge until
315 a saturation point is reached. At high thresholds,
316 the image becomes nearly black, and topological
317 features disappear. This dynamic is further illus-
318 trated in Figure 3, which shows the effect of different
319 threshold values on the R, G, and B channels of an
320 image. The transformation of pixel intensities across
321 channels explains the bell-shaped nature of the Betti
322 curve. It is noteworthy to mention that the RGB
323 color space in the dataset is the default mode, which
324 is aligned with the staining of the tissues. [50] pro-
325 vides a comprehensive overview of different color
326 spaces in digital pathology.

327 Since our dataset consists of RGB images, we
328 applied the described cubical complex pipeline sep-
329 arately to each channel. Consequently, each image
330 yields three Betti curves—one for each channel.

331 These topological descriptors are then concatenated
332 together and forwarded to the CBMIR pipeline for
333 downstream analysis and similarity-based image
334 search.

335 Betti curves are naturally calculated on a non-
336 uniform filtration scale. For instance, let us
337 assume we have n loops represented as $\beta_1 =$
338 $[(b_1, d_1), (b_2, d_2), \dots, (b_n, d_n)]$. For homogeneous
339 treatment across datasets, we uniformly select $R = i$
340 filtration points within the range between the mini-
341 mum birth value and the maximum death value, de-
342 noted as $X = [R_1, R_2, \dots, R_i]$. A loop $\beta_1(n)$ is
343 considered alive at a filtration point X_i if $b_n \leq X_i \leq d_n$.
344 Thus, we introduce a Betti curve resolution (R),
345 which determines the granularity at which topologi-
346 cal features are detected such that higher resolutions
347 preserve fine-grained structures but may capture
348 noisy artifacts, while lower resolutions emphasize
349 large-scale features at the cost of missing subtle pat-
350 terns. The appropriate resolution for a given image
351 set is a modeling choice that balances sensitivity
352 and robustness in tissue analysis.

4 Experiments 353

354 In this paper, we focus on the application of TDA
355 in digital pathology and analyze the topological pat-
356 terns of the histopathological images. To do so, our
357 method of choice is the Betti curves of the cubical
358 complexes. The cubic complex from the Gudhi Li-
359 brary² works with grayscale images. However, in
360 digital pathology and working with WSIs, color plays
361 a significant role that cannot be discarded from the
362 experiments [40]. To address this, we apply cubical
363 complex persistence separately to each RGB chan-
364 nel, treating them as independent grayscale images.

²<https://gudhi.inria.fr/cubicalcomplex/>




Figure 3. shows three channels of an RGB image under different th values. These values were defined as a sub-level filtration in $[0-1]$ for the normalized images. Each channel goes through the same th values to illustrate the th impacts on channels.

365 The resulting Betti values from each channel are
 366 then concatenated to form a comprehensive topo-
 367 logical descriptor for each image. With Betti curves
 368 sampled uniformly in R resolution, for each image
 369 we have $3 \times R$ features. As mentioned above, finding
 370 the optimum value for the resolution is challenging
 371 and has a direct impact on the final results. As the
 372 resolution R increases, more fine-grained topological
 373 details are captured, though at the cost of longer
 374 computation times. In the case of BreakHis data
 375 set, $R = 200$ offers a strong balance between per-
 376 formance and efficiency, yielding the best accuracy
 377 with a much lower computational burden, followed
 378 by [45]. So, for the following experiments in this
 379 paper, we considered $R = 200$. Therefore, a descrip-
 380 tor representing each image is constructed from the
 381 features derived from β_1 for the CBMIR framework.
 382 In the next step, the extracted features for the en-
 383 tire dataset are saved. Subsequently, we compute
 384 the features of the test set, which was previously
 385 separated from the training set. With 600 features
 386 extracted per image, we proceed to the comparison
 387 phase. Specifically, we use the Euclidean distance
 388 metric to identify the top- K most similar patches
 389 from the training set for each query image. After
 390 ranking the images based on their smallest distances,
 391 the top matching patches are retrieved and visual-
 392 ized along with their corresponding labels to assist
 393 pathologists in analysis.

5 Results and discussion

395 When evaluating the performance of CBMIR meth-
 396 ods, several important aspects must be considered,
 397 including training time, accuracy, and ease of train-
 398 ing. In this section, we compare the performance of
 399 THIR with other methods across all these dimen-
 400 sions.

5.1 Accuracy Comparison

401 One challenge in this study is the scarcity of compa-
 402 rable methods evaluated on the same dataset under
 403 identical conditions (i.e., same K value and magni-
 404 fication). To address this, we provide Table 2 and
 405 some more figures, such as Figure 4, to provide a
 406 comprehensive overview of the results in this study.

407 summarizing recent studies on BreakHis, indicat-
 408 ing their magnification and the K value considered.
 409 This enables clearer benchmarking of THIR across
 410 multiple settings. The evaluation is performed at
 411 multiple values of K , focusing on key performance
 412 metrics: accuracy, recall, precision, and F1-score.

413 In published studies on CBMIR, the value of K
 414 varies across works. Additionally, different magni-
 415 fications of the BreakHis dataset have been con-
 416 sidered, and only a few studies have applied their
 417 methods to all magnifications ($40\times$, $100\times$, $200\times$,
 418 and $400\times$). These variations make direct compari-
 419 son challenging. To address this, we applied THIR
 420 across all magnifications of the data set using the
 421 most commonly used values of K . Following the
 422 evaluation protocol of [32, 51, 52], we conducted a
 423 fair comparison with several state-of-the-art meth-
 424 ods.

425 At $400\times$ magnification, THIR achieves an accu-
 426 racy of 0.98 at the top-5, outperforming both super-
 427 vised and unsupervised methods. Notably, it delivers
 428 an accuracy of 0.98, which is 18% higher than Breast-
 429 twins (0.69), a fully supervised method utilizing a
 430 Siamese network for similarity learning. In [14], Fed-
 431 CBMIR was designed to generalize across all mag-
 432 nifications and was evaluated at $400\times$. Although
 433 FedCBMIR aims to enhance performance on unseen
 434 data, THIR consistently delivers better retrieval
 435 accuracy and precision, even without any learning
 436 phase or hyperparameter tuning. This trend contin-
 437 ues at $200\times$ magnification with $K = 5$, where THIR
 438 achieves a precision of 0.99, surpassing FedCBMIR
 439 by 10% (FedCBMIR precision = 0.89). Addition-
 440 ally, THIR demonstrates stronger performance than
 441 other baselines like CNN-based AE (0.93 precision)
 442 and MCCH (0.89 precision).

443 In Table 2, we include three methods from [53],
 444 for which the value of K used in top- K retrieval
 445 is not explicitly defined in the original paper. The
 446 only instance where a specific K value is mentioned
 447 appears in the caption of a qualitative figure, where
 448 they state $K = 5$. However, this information is
 449 not provided for the quantitative results reported
 450 in the tables. Therefore, due to the lack of clarity
 451 regarding the retrieval threshold, we marked the K
 452 value as "Not-defined" in our comparison table. The
 453 paper [53] focuses on hashing-based methods and
 454 reports retrieval performance at various code lengths:
 455 16, 32, and 64 bits. To include these results in our
 456 comparison, we selected the highest performance
 457 across all bit lengths. For example, the DCMMH
 458

459 method achieves its best performance (0.95) at 32
460 bits for $40\times$ magnification, while DPSH reaches its
461 highest at 16 bits for the same magnification. To
462 ensure a fair comparison with our proposed method,
463 we report only the top-performing results from each
464 approach, regardless of the bit length used.

465 Across all magnifications and for both $K = 3$,
466 $K = 5$, and $K = \text{Not defined}$, THIR maintains
467 high and stable performance. It consistently outper-
468 forms CNN-based AE, MCCH, and several hashing-
469 based methods (e.g., HashNet, IDHN, DTQ), as well
470 as state-of-the-art frameworks such as FedCBMIR
471 and VTHC. These results highlight the robustness
472 and effectiveness of TDA in CBMIR, especially com-
473 pared to DL-based methods that require substantial
474 training time and parameter optimization.

475 Figure 4 shows four random examples of retrieval
476 results with corresponding labels. Incorrect retri-
477 eivals (where the retrieved label differs from the
478 query label) are outlined in red. This visual analysis
479 enables qualitative evaluation of system performance
480 and highlights the capability of topological features
481 to capture structural similarities in histopathologi-
482 cal breast cancer images. Beyond label checking,
483 pathologists can also examine structural patterns
484 between queries and retrieved results based on their
485 expertise.

486 Figure 5 contains four panels, each showing the
487 Betti curves of four randomly selected images. Differ-
488 ent colored lines represent different images. A guide
489 bar on top of each panel indicates the corresponding
490 class label of each image related to its curve. The
491 x-axis displays the filtration steps, while the y-axis
492 shows the Betti values. The x-axis ranges from 0
493 to 600, illustrating that 600 features were extracted
494 from the Betti values of each RGB image. As men-
495 tioned earlier, we computed the Betti values for each
496 channel separately using the cubical complex and

497 then concatenated them to obtain a representative
498 feature vector for the entire image. Thus, the x-axis
499 can be interpreted in intervals of 200: the interval
500 $[0-200]$ represents the Betti values for the red chan-
501 nel, $[200-400]$ for green, and $[400-600]$ for blue. This
502 figure demonstrates how the trend of Betti curves
503 behaves across channels for images from the same or
504 different classes. For instance, in the top-left panel,
505 all images belong to the same class (Benign), and
506 their Betti curves follow a similar trend, indicating
507 topological similarity. In contrast, in the top-right
508 panel, the red line deviates noticeably, suggesting a
509 different topological pattern compared to the other
510 three images. The remaining three (blue, orange,
511 and green lines) follow a similar trend, which reflects
512 their similarity in class labels. In the bottom panels,
513 two additional examples show that two images have
514 similar Betti curves and share the same labels, while
515 the other two follow distinct trends and have differ-
516 ent labels. These characteristic patterns suggest
517 that Betti curves encode discriminative information
518 suitable for unsupervised CBMIR.

519 Furthermore, Figure 6 demonstrates the retrieval
520 of four random queries based on Betti values. In
521 each retrieval panel, four Betti curves are shown:
522 the red line represents the query, while blue, orange,
523 and green lines represent the top-3 retrieved images.
524 The class labels are indicated above each image,
525 and a guide bar links the curve colors to the im-
526 ages. In some cases, such as the top-left panel, the
527 retrieved image shares similar topological features
528 with the query but has a different label, suggesting
529 intraobserver variability. In other cases, such as the
530 top-right panel, all retrieved images share the same
531 label with the query, reinforcing the robustness of
532 Betti features. Such visualization suggests that im-
533 ages with similar Betti curves may share underlying
534 histopathological characteristics, even when labeled
535 differently. This highlights the potential of CBMIR
536 for digital pathology applications and demonstrates
537 advantages over traditional Computer-Aided Diag-
538 nosis (CAD) tools.

539 As an indirect comparison between the THIR
540 result and state-of-the-art classifiers on the same
541 images, Table 3 provides comprehensive information
542 regarding the accuracy of classifiers on the BreaKHis
543 data set at $40X$. TopOC-1 and TopOC-CNN in [45]
544 obtained 89% and 93% accuracy, while THIR was
545 successful in retrieving images with the same label
546 with 98% accuracy. However, TopOC-1 and TopOC-
547 CNN needed some training time and hyperparameter
548 tuning.




Figure 4. shows four random queries and their similar patches. Each row represents a query image (leftmost) followed by its top-3 retrieved images. The true class labels are shown as Query Label and Retrieved Label. Misclassified retrievals are outlined in red.

5.2 How Fast, How Simple: Training, Searching, and Using the Model

Several recent CBMIR studies based on CAE have reported training times of approximately 9.33 hours

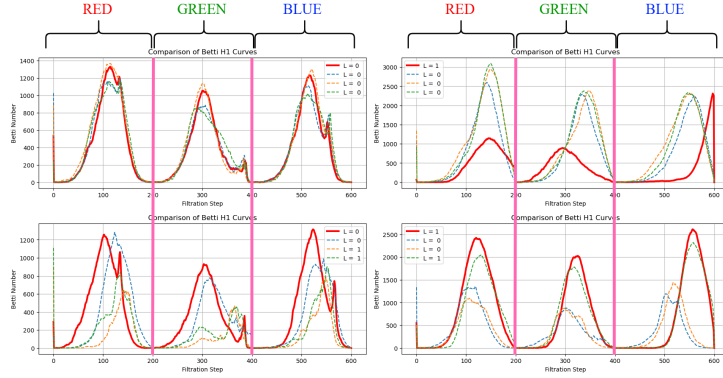


Figure 5. shows four panels, each with four concatenated Betti curves for four random images. The label of each image is represented in a small bar on top of each panel. $L = 0$ and $L = 1$ mean "Benign" and "Malignant" cases, respectively. This explains how images with the same cancer grade have similar Betti curves. The y-axis illustrates the number of loops at each filtration step, and the x-axis represents the filtration step for each channel. $R = 200$ yields 200 filter steps for each channel, which means 600 filter steps in total for an RGB image.

Table 2. Performance comparison of THIR with recent methods at various K values on BreakHis across all magnifications. The best results per metric are bolded.

Magnification	Method	K	Accuracy	Recall	Precision	F1-score
40×	THIR	3	0.95	0.97	0.96	0.97
	THIR	5	0.98	0.97	0.96	0.95
	FedCBMIR [14]	5	0.97	-	0.96	0.98
	CBMIR [14]		0.95	-	0.93	0.96
	MCCH [54]		-	-	0.94	-
	CNN-based AE [52]		-	0.77	0.95	-
	HSDH [55]	400	-	-	0.99	-
	DTQ [55]		-	-	0.91	-
	ATH [55]		-	-	0.89	-
	IDHN [55]		-	-	0.95	-
HashNet [55]	-	-	0.91	-		
VTHC [56]	6	-	-	0.98	-	
DCMMH [53]	Not defined	-	-	0.96	-	
DPSH [53]		-	-	0.95	-	
ADSH [53]		-	-	0.95	-	
100×	THIR	3	0.97	0.98	0.98	0.98
	THIR	5	0.99	0.99	0.99	0.99
	FedCBMIR [14]	5	0.94	-	0.92	0.96
	CBMIR [14]		0.90	-	0.88	0.94
	MCCH [54]		-	-	0.92	-
	CNN-based AE [52]		-	0.49	0.77	-
	VTHC [56]	6	-	-	0.99	-
	DCMMH [53]	Not defined	-	-	0.95	-
DPSH [53]	-		-	0.95	-	
ADSH [53]	-		-	0.94	-	
200×	THIR	3	0.97	0.98	0.98	0.98
	THIR	5	0.99	0.99	0.99	0.99
	FedCBMIR [14]	5	0.92	-	0.89	0.94
	CBMIR [14]		0.89	-	0.87	0.93
	MCCH [54]		-	-	0.91	-
	CNN-based AE [52]		-	0.76	0.92	-
	VTHC [56]	6	-	-	0.98	-
	DCMMH [53]	Not defined	-	-	0.97	-
DPSH [53]	-		-	0.96	-	
ADSH [53]	-		-	0.95	-	
400×	THIR	3	0.95	0.97	0.95	0.97
	THIR	5	0.98	0.99	0.98	0.99
	FedCBMIR [14]	5	0.96	-	0.94	0.97
	UCBMIR [30]		-	0.79	0.91	-
	Breast-twins [54]		0.69	0.82	0.91	0.81
	MCCH [31]		-	-	0.89	-
	CNN-based AE [52]	-	0.69	0.93	-	
	VTHC [56]	6	-	-	0.99	-
DCMMH [53]	Not defined	-	-	0.96	-	
DPSH [53]		-	-	0.95	-	
ADSH [53]		-	-	0.95	-	

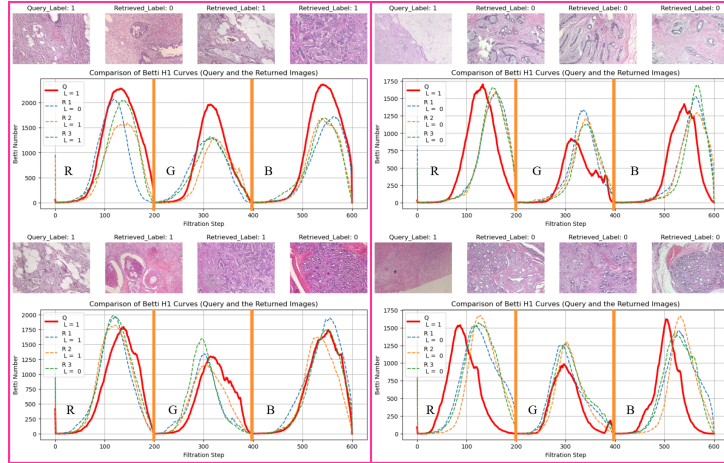


Figure 6. demonstrates the retrieval of four random queries using Betti values. There are four panels of Betti curves with the corresponding queries and returned images. For each panel of Betti curves, a query image (top left) is compared to a set of retrieved images using the Euclidean distance function. Each curve in the panel reflects topological features over the filtration steps. The alignment of Betti curve patterns with the query supports the effectiveness of PH for topology-aware image retrieval. The label of each image is represented in a small bar on top of each panel. $L = 0$ and $L = 1$ mean "Benign" and "Malignant" images, respectively.

553 and 6.59 hours on the BreakeHis dataset [14] at $40\times$
 554 magnification, using an NVIDIA Tesla T4 GPU. Al-
 555 though FedCBMIR reduces training time compared
 556 to traditional CBMIR, it still requires a substan-
 557 tial training duration. In DL-based methods, once
 558 the models are trained, features are subsequently
 559 extracted for CBMIR tasks; however, the feature
 560 extraction time is often not reported in these studies.
 561 In contrast, our method bypasses the lengthy train-
 562 ing phase entirely. THIR efficiently processes the
 563 entire dataset by performing the filtration, calculat-
 564 ing the Betti values β_1 , and extracting 600 features
 565 per image. This comprehensive feature extraction
 566 for all images in the dataset requires merely about
 567 20 minutes, which is significantly faster than the
 568 combined training and feature extraction time of
 569 existing deep learning approaches.

Table 3. Classification accuracy of THIR compared to state-of-the-art methods on the BreakeHis dataset at $400\times$ magnification, with $K = 5$ for THIR.

Method	Accuracy
THIR	0.98
DenseNet201 [57]	0.95
IDSNet [58]	0.94
VGG16 [59]	0.94
Resnet [60]	0.94
DenseNet201 [61]	0.89
BkCapsNet [30]	0.88
CapsNet [62]	0.88
BkNet [63]	0.84
MobileNet [64]	0.84
AlexNet [65]	0.81

6 Conclusion and future work 570

This paper introduces THIR, an unsupervised, 571
 training-free, and interpretable framework for 572
 Content-Based Medical Image Retrieval (CBMIR) 573
 based on topological data analysis. Using cubical 574
 persistence and Betti values for loops (β_1), THIR 575
 extracts robust topological fingerprints from raw 576
 RGB histopathological images without annotations, 577
 GPU acceleration, or hyperparameter tuning. 578

Experiments on the BreakeHis dataset show that 579
 THIR consistently outperforms supervised and un- 580
 supervised baselines across all magnifications. At 581
 $400\times$, it improves accuracy by 31% and precision 582
 by 18% over Breast-twins, a Siamese-network-based 583
 model. At $200\times$, it achieves a 10% gain in precision 584
 over FedCBMIR. Similar margins are observed 585
 at lower magnifications, with precision up to 0.99 586
 at $100\times$, surpassing all compared methods. These 587
 results confirm THIR as a strong alternative for CB- 588
 MIR in label-scarce or resource-constrained settings. 589

THIR extracts 600 topological features per image 590
 in under 20 minutes on a standard CPU, offering 591
 an efficient and scalable solution compared to deep 592
 learning models requiring extensive training. Its 593
 performance, efficiency, and interpretability make 594
 it suitable for clinical applications and diagnostic 595
 support. 596

Future work will explore higher-dimensional fea- 597
 tures (e.g., persistence images, landscapes), exten- 598
 sion to multi-channel and multi-organ datasets, alter- 599
 native color spaces beyond RGB, and generalization 600
 to multi-class and whole-slide image retrieval tasks. 601

602 **References**

- 603 [1] F. Bray, M. Laversanne, H. Sung, J. Ferlay,
604 R. L. Siegel, I. Soerjomataram, and A. Jemal.
605 “Global cancer statistics 2022: GLOBOCAN es-
606 timates of incidence and mortality worldwide
607 for 36 cancers in 185 countries”. In: *CA: a can-
608 cer journal for clinicians* 74.3 (2024), pp. 229–
609 263.
- 610 [2] M. Arnold, E. Morgan, H. Rungay, A. Mafra,
611 D. Singh, M. Laversanne, J. Vignat, J. R.
612 Gralow, F. Cardoso, S. Siesling, et al. “Cur-
613 rent and future burden of breast cancer: Global
614 statistics for 2020 and 2040”. In: *The Breast*
615 66 (2022), pp. 15–23.
- 616 [3] N. Kanwal, F. Khoraminia, U. Kiraz, A.
617 Mosquera-Zamudio, C. Monteagudo, E. A.
618 Janssen, T. C. Zuiverloon, C. Rong, and K.
619 Engan. “Equipping computational pathology
620 systems with artifact processing pipelines: a
621 showcase for computation and performance
622 trade-offs”. In: *BMC Medical Informatics and
623 Decision Making* 24.1 (2024), p. 288.
- 624 [4] A. M. Khan, K. Sirinukunwattana, and N.
625 Rajpoot. “A global covariance descriptor for
626 nuclear atypia scoring in breast histopathology
627 images”. In: *IEEE journal of biomedical and
628 health informatics* 19.5 (2015), pp. 1637–1647.
- 629 [5] H. Singh, S. Sethi, M. Raber, and L. A. Pe-
630 tersen. “Errors in cancer diagnosis: current
631 understanding and future directions”. In: *Jour-
632 nal of clinical oncology* 25.31 (2007), pp. 5009–
633 5018.
- 634 [6] O. Kostopoulou, B. C. Delaney, and C. W.
635 Munro. “Diagnostic difficulty and error in pri-
636 mary care—a systematic review”. In: *Family
637 practice* 25.6 (2008), pp. 400–413.
- 638 [7] M. R. Abbasniya, S. A. Sheikholeslamzadeh,
639 H. Nasiri, and S. Emami. “Classification of
640 breast tumors based on histopathology images
641 using deep features and ensemble of gradient
642 boosting methods”. In: *Computers and Elec-
643 trical Engineering* 103 (2022), p. 108382.
- 644 [8] C. Chen, M. Y. Lu, D. F. Williamson, T. Y.
645 Chen, A. J. Schaumberg, and F. Mahmood.
646 “Fast and scalable search of whole-slide images
647 via self-supervised deep learning”. In: *Nature
648 Biomedical Engineering* 6.12 (2022), pp. 1420–
649 1434.
- 650 [9] R. C. Veltkamp and M. Tanase. *Content-based
651 image retrieval systems: A survey*. Tech. rep.
652 Technical Report UU-CS-2000-34, Dept. of
653 Computing Science, Utrecht University, 2000.
- [10] D. Komura and S. Ishikawa. “Machine learning
654 methods for histopathological image analysis”.
655 In: *Computational and structural biotechnology
656 journal* 16 (2018), pp. 34–42. 657
- [11] S. Kalra, H. R. Tizhoosh, S. Shah, C. Choi,
658 S. Damaskinos, A. Safarpour, S. Shafiei, M.
659 Babaie, P. Diamandis, C. J. Campbell, et
660 al. “Pan-cancer diagnostic consensus through
661 searching archival histopathology images us-
662 ing artificial intelligence”. In: *NPJ digital
663 medicine* 3.1 (2020), p. 31. 664
- [12] A. Kumar, F. Nette, K. Klein, M. Fulham, and
665 J. Kim. “A visual analytics approach using the
666 exploration of multidimensional feature spaces
667 for content-based medical image retrieval”. In:
668 *IEEE journal of biomedical and health infor-
669 matics* 19.5 (2014), pp. 1734–1746. 670
- [13] H. R. Tizhoosh and L. Pantanowitz. “On im-
671 age search in histopathology”. In: *Journal of
672 Pathology Informatics* (2024), p. 100375. 673
- [14] Z. Tabatabaei, Y. Wang, A. Colomer, J. Oliver
674 Moll, Z. Zhao, and V. Naranjo. “Wwfedcb-
675 mir: World-wide federated content-based med-
676 ical image retrieval”. In: *Bioengineering* 10.10
677 (2023), p. 1144. 678
- [15] Y. Ma, Z. Jiang, H. Zhang, F. Xie, Y. Zheng,
679 H. Shi, and Y. Zhao. “Breast histopathological
680 image retrieval based on latent dirichlet allo-
681 cation”. In: *IEEE journal of biomedical and
682 health informatics* 21.4 (2016), pp. 1114–1123. 683
- [16] T. Bhaskar, Y. Ramadevi, P. N. Kavitha, and
684 P. Sravan. “MCBIR: Deep Learning based
685 Framework for Efficient Content Based Im-
686 age Retrieval System of Medical Images”. In:
687 (). 688
- [17] C. Liu, X. Ma, H. Zhang, S. Xie, and D. Yu.
689 “Dynamic Neuropsychological Approach for
690 Multi-Quality Image Assessment Using Grey-
691 Topological Data Analysis”. In: *IEEE Access*
692 (2024). 693
- [18] M. Tian, M. Su, X. Xiao, S. Yi, Z. Hua, and
694 Y. Zhang. “High-precision privacy-protected
695 image retrieval based on multi-feature fusion”.
696 In: *Knowledge-Based Systems* 315 (2025),
697 p. 113243. 698
- [19] U. Sharif, Z. Mehmood, T. Mahmood, M. A.
699 Javid, A. Rehman, and T. Saba. “Scene analy-
700 sis and search using local features and support
701 vector machine for effective content-based im-
702 age retrieval”. In: *Artificial Intelligence Review*
703 52 (2019), pp. 901–925. 704

- [20] A. Sarwar, Z. Mehmood, T. Saba, K. A. Qazi, A. Adnan, and H. Jamal. "A novel method for content-based image retrieval to improve the effectiveness of the bag-of-words model using a support vector machine". In: *Journal of Information Science* 45.1 (2019), pp. 117–135.
- [21] Z. Mehmood, F. Abbas, T. Mahmood, M. A. Javid, A. Rehman, and T. Nawaz. "Content-based image retrieval based on visual words fusion versus features fusion of local and global features". In: *Arabian Journal for Science and Engineering* 43.12 (2018), pp. 7265–7284.
- [22] C. S. Won, D. K. Park, and S.-J. Park. "Efficient use of MPEG-7 edge histogram descriptor". In: *ETRI journal* 24.1 (2002), pp. 23–30.
- [23] B. S. Manjunath, J.-R. Ohm, V. V. Vasudevan, and A. Yamada. "Color and texture descriptors". In: *IEEE Transactions on circuits and systems for video technology* 11.6 (2001), pp. 703–715.
- [24] M. Owais, M. Arsalan, J. Choi, and K. R. Park. "Effective diagnosis and treatment through content-based medical image retrieval (CB-MIR) by using artificial intelligence". In: *Journal of clinical medicine* 8.4 (2019), p. 462.
- [25] A. Widmer, R. Schaer, D. Markonis, and H. Müller. "Gesture interaction for content-based medical image retrieval". In: *Proceedings of international conference on multimedia retrieval*. 2014, pp. 503–506.
- [26] K. Karthik and S. S. Kamath. "A deep neural network model for content-based medical image retrieval with multi-view classification". In: *The Visual Computer* 37.7 (2021), pp. 1837–1850.
- [27] V. T. H. Tuyet, N. T. Binh, N. K. Quoc, and A. Khare. "Content based medical image retrieval based on salient regions combined with deep learning". In: *Mobile Networks and Applications* 26.3 (2021), pp. 1300–1310.
- [28] S. Agrawal, A. Chowdhary, S. Agarwala, V. Mayya, and S. Kamath S. "Content-based medical image retrieval system for lung diseases using deep CNNs". In: *International Journal of Information Technology* 14.7 (2022), pp. 3619–3627.
- [29] S. R. Dubey. "A decade survey of content based image retrieval using deep learning". In: *IEEE Transactions on Circuits and Systems for Video Technology* 32.5 (2021), pp. 2687–2704.
- [30] Z. Tabatabaei, A. Colomer, J. O. Moll, and V. Naranjo. "Toward More Transparent and Accurate Cancer Diagnosis With an Unsupervised CAE Approach". In: *IEEE Access* 11 (2023), pp. 143387–143401. DOI: [10.1109/ACCESS.2023.3343845](https://doi.org/10.1109/ACCESS.2023.3343845).
- [31] Z. Tabatabaei, A. Colomer, J. O. Moll, and V. Naranjo. "Siamese Content-based Search Engine for a More Transparent Skin and Breast Cancer Diagnosis through Histological Imaging". In: *arXiv preprint arXiv:2401.08272* (2024).
- [32] N. Hegde, J. D. Hipp, Y. Liu, M. Emmert-Buck, E. Reif, D. Smilkov, M. Terry, C. J. Cai, M. B. Amin, C. H. Mermel, et al. "Similar image search for histopathology: SMILY". In: *NPJ digital medicine* 2.1 (2019), p. 56.
- [33] S. Kalra, H. R. Tizhoosh, C. Choi, S. Shah, P. Diamandis, C. J. Campbell, and L. Pantanowitz. "Yottixel—an image search engine for large archives of histopathology whole slide images". In: *Medical Image Analysis* 65 (2020), p. 101757.
- [34] X. Wang, Y. Du, S. Yang, J. Zhang, M. Wang, J. Zhang, W. Yang, J. Huang, and X. Han. "RetCCL: Clustering-guided contrastive learning for whole-slide image retrieval". In: *Medical image analysis* 83 (2023), p. 102645.
- [35] Y. Ozen, S. Aksoy, K. Kösemehmetoğlu, S. Önder, and A. Üner. "Self-supervised learning with graph neural networks for region of interest retrieval in histopathology". In: *2020 25th International conference on pattern recognition (ICPR)*. IEEE. 2021, pp. 6329–6334.
- [36] Y. Zheng, Z. Jiang, H. Zhang, F. Xie, Y. Ma, H. Shi, and Y. Zhao. "Size-scalable content-based histopathological image retrieval from database that consists of WSIs". In: *IEEE journal of biomedical and health informatics* 22.4 (2017), pp. 1278–1287.
- [37] Y. Skaf and R. Laubenbacher. "Topological data analysis in biomedicine: A review". In: *Journal of Biomedical Informatics* 130 (2022), p. 104082.
- [38] A. H. Abdulaal, M. Valizadeh, R. A. Yassin, M. C. Amirani, A. S. Shah, B. M. Albaker, and A. S. M. Mustaf. "Hybrid CNN and RNN Model for Histopathological Sub-Image Classification in Breast Cancer Analysis Using Self-Learning". In: *Journal of Engineering and Sustainable Development* 29.3 (2025), pp. 310–320.

- [39] N. Kanwal, F. Pérez-Bueno, A. Schmidt, K. Engan, and R. Molina. “The devil is in the details: Whole slide image acquisition and processing for artifacts detection, color variation, and data augmentation: A review”. In: *Ieee Access* 10 (2022), pp. 58821–58844.
- [40] W. Al Noumah, A. Jafar, and K. Al Joumaa. “Using parallel pre-trained types of DCNN model to predict breast cancer with color normalization”. In: *BMC research notes* 15 (2022), pp. 1–6.
- [41] F. Wang, S. Kapse, S. Liu, P. Prasanna, and C. Chen. “TopoTxR: a topological biomarker for predicting treatment response in breast cancer”. In: *International Conference on Information Processing in Medical Imaging*. Springer. 2021, pp. 386–397.
- [42] L. Crawford, A. Monod, A. X. Chen, S. Mukherjee, and R. Rabadán. “Predicting clinical outcomes in glioblastoma: an application of topological and functional data analysis”. In: *Journal of the American Statistical Association* 115.531 (2020), pp. 1139–1150.
- [43] J. Cui. “Extended persistence and duality in cubical complexes”. In: (2024).
- [44] D. Strömbom. *Persistent homology in the cubical setting: theory, implementations and applications*. 2007.
- [45] S. Fatema, B. Nuwagira, S. Chakraborty, R. Gedik, and B. Coskunuzer. “TopOC: Topological Deep Learning for Ovarian and Breast Cancer Diagnosis”. In: *International Workshop on Topology and Graph-Informed Imaging Informatics*. Springer. 2024, pp. 22–32.
- [46] N. Otter, M. A. Porter, U. Tillmann, P. Grindrod, and H. A. Harrington. “A roadmap for the computation of persistent homology”. In: *EPJ Data Science* 6 (2017), pp. 1–38.
- [47] R. Rabadán and A. J. Blumberg. *Topological data analysis for genomics and evolution: topology in biology*. Cambridge University Press, 2019.
- [48] M. Dugast, G. Bouleux, O. Mory, and E. Marcon. “Improving health care management through persistent homology of time-varying variability of emergency department patient flow”. In: *IEEE journal of biomedical and health informatics* 23.5 (2018), pp. 2174–2181.
- [49] P. Lawson, J. Schupbach, B. T. Fasy, and J. W. Sheppard. “Persistent homology for the automatic classification of prostate cancer aggressiveness in histopathology images”. In: *Medical Imaging 2019: Digital Pathology*. Vol. 10956. SPIE. 2019, pp. 72–85.
- [50] R. Velastegui and M. Pedersen. “The impact of using different color spaces in histological image classification using convolutional neural networks”. In: *2021 9th European Workshop on Visual Information Processing (EUVIP)*. IEEE. 2021, pp. 1–6.
- [51] S. Murala and Q. J. Wu. “Local mesh patterns versus local binary patterns: biomedical image indexing and retrieval”. In: *IEEE journal of biomedical and health informatics* 18.3 (2013), pp. 929–938.
- [52] A. E. Minarno, K. M. Ghufro, T. S. Sabrila, L. Husniah, and F. D. S. Sumadi. “Cnn based autoencoder application in breast cancer image retrieval”. In: *2021 International Seminar on Intelligent Technology and Its Applications (ISITIA)*. IEEE. 2021, pp. 29–34.
- [53] Y. Gu and J. Yang. “Densely-connected multi-magnification hashing for histopathological image retrieval”. In: *IEEE journal of biomedical and health informatics* 23.4 (2018), pp. 1683–1691.
- [54] Y. Gu and J. Yang. “Multi-level magnification correlation hashing for scalable histopathological image retrieval”. In: *Neurocomputing* 351 (2019), pp. 134–145.
- [55] S. M. Alizadeh, M. S. Helfroush, and H. Müller. “A novel Siamese deep hashing model for histopathology image retrieval”. In: *Expert Systems with Applications* 225 (2023), p. 120169.
- [56] M. Kumar, R. Singh, and P. Mukherjee. “VTHSC-MIR: Vision Transformer Hashing with Supervised Contrastive learning based medical image retrieval”. In: *Pattern Recognition Letters* 184 (2024), pp. 28–36.
- [57] F. Taheri and K. Rahbar. “Enhancing breast cancer diagnosis: transfer learning on DenseNet with neural hashing for histopathology fine-grained image classification”. In: *Medical & Biological Engineering & Computing* (2025), pp. 1–15.
- [58] X. Li, X. Shen, Y. Zhou, X. Wang, and T.-Q. Li. “Classification of breast cancer histopathological images using interleaved DenseNet with SENet (IDSNet)”. In: *PloS one* 15.5 (2020), e0232127.
- [59] Y. Liang and Z. Meng. “Brea-net: An interpretable dual-attention network for imbalanced breast cancer classification”. In: *IEEE Access* 11 (2023), pp. 100508–100517.
- [60] F. B. Ashraf, S. M. Alam, and S. M. Sakib. “Enhancing breast cancer classification via histopathological image analysis: Leveraging self-supervised contrastive learning and transfer learning”. In: *Heliyon* 10.2 (2024).

- 920 [61] A. Maleki, M. Raahemi, and H. Nasiri. “Breast
921 cancer diagnosis from histopathology images
922 using deep neural network and XGBoost”. In:
923 *Biomedical Signal Processing and Control* 86
924 (2023), p. 105152.
- 925 [62] S. Sabour, N. Frosst, and G. E. Hinton. “Dy-
926 namic routing between capsules”. In: *Advances*
927 *in neural information processing systems* 30
928 (2017).
- 929 [63] P. Wang, J. Wang, Y. Li, P. Li, L. Li, and M.
930 Jiang. “Automatic classification of breast can-
931 cer histopathological images based on deep fea-
932 ture fusion and enhanced routing”. In: *Biomed-*
933 *ical Signal Processing and Control* 65 (2021),
934 p. 102341.
- 935 [64] E. O. Simonyan, J. A. Badejo, and J. S. Weijin.
936 “Histopathological breast cancer classification
937 using CNN”. In: *Materials Today: Proceedings*
938 105 (2024), pp. 268–275.
- 939 [65] F. A. Spanhol, L. S. Oliveira, C. Petitjean,
940 and L. Heutte. “Breast cancer histopatholog-
941 ical image classification using convolutional
942 neural networks”. In: *2016 international joint*
943 *conference on neural networks (IJCNN)*. IEEE.
944 2016, pp. 2560–2567.