To appear in the Proceedings of the 19th IEEE Pacific Visualization Conference (PacificVis 2026), Sydney, Australia (Conference Track).

arXiv:2602.13402v1 [cs.HC] 13 Feb 2026

InfoCIR: Multimedia Analysis for Composed Image Retrieval

Ioannis Dravilas*, †, ‡ Ioannis Kapetangeorgis*, ‡ Anastasios Latsoudis*, ‡ Conor McCarthy§
Gonçalo Marcelino§ Marcel Worring§

University of Amsterdam, The Netherlands

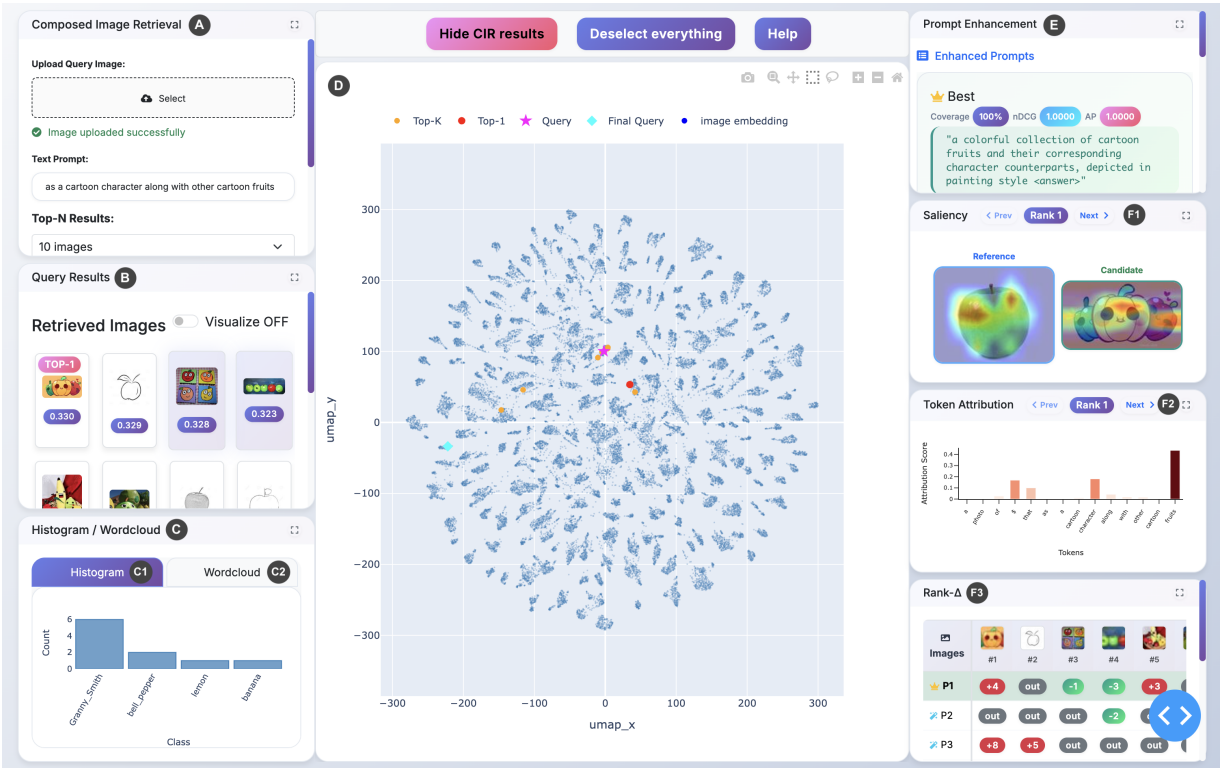

Figure 1: The user interface of our system consists of six main panels. (A) The Composed Image Retrieval Panel allows users to input an image and a text prompt to initiate the retrieval process and establish a baseline search. (B) The Query Results Panel displays the top-k images which are listed by similarity; any image can be clicked to mark it as an ideal target for prompt refinement. (C) The Histogram / Word cloud Panel includes the class-frequency histogram (C_1) and a word cloud (C_2), summarizing labels within the current top-k. (D) The central Embedding View shows a 2D UMAP projection of the dataset, highlighting the reference image, the composed query embedding, and the top-k results. (E) The Prompt Enhancement Panel proposes alternative prompts conditioned on the user-selected ideal target images, using an LLM, and retrieval metrics. (F) The Explanation Panel visualizes model attribution using a saliency map (F_1), a token attribution bar chart (F_2) and a Rank-Δ heatmap (F_3).

Abstract

Composed Image Retrieval (CIR) allows users to search for images by combining a reference image with a text prompt that describes desired modifications. While vision-language models like CLIP have popularized this task by embedding multiple modalities into a joint space, developers still lack tools that reveal how these multimodal prompts interact with embedding spaces and why small wording changes can dramatically alter the results. We present InfoCIR, a visual analytics system that closes this gap by coupling retrieval, explainability, and prompt engineering in a single, interactive dashboard. InfoCIR integrates a state-of-the-art CIR back-end (SEARLE [1]) with a six-panel interface that (i) lets users compose image + text queries, (ii) projects the top-k results into a low-dimensional space using Uniform Manifold Approximation and Projection (UMAP) for spatial reasoning, (iii) overlays similarity-based saliency maps and gradient-derived token-attribution bars for local explanation, and (iv) employs an LLM-powered prompt enhancer that generates counterfactual variants and visualizes how these changes affect the ranking of user-selected target images. A modular architecture built on Plotly-Dash allows new models, datasets, and attribution methods to be plugged in with minimal

*Contributed equally to this work.
†Corresponding author: Ioannis Dravilas
‡e-mail: {ioannis.dravilas, ioannis.kapetangeorgis, anastasios.latsoudis}@student.uva.nl
§e-mail: {c.t.mccarthy, g.barretoferreiramarcelino, m.worring}@uva.nl
© 2026 IEEE. Personal use of this material is permitted. Permission from IEEE must be obtained for all other uses, in any current or future media, including reprinting/republishing this material for advertising or promotional purposes, creating new collective works, for resale or redistribution to servers or lists, or reuse of any copyrighted component of this work in other works.

effort. We argue that InfoCIR helps diagnose retrieval failures, guides prompt enhancement, and accelerates insight generation during model development. All source code allowing for a reproducible demo is available at https://github.com/giannhskp/InfoCIR.

Index Terms: Composed Image Retrieval, Prompt Enhancement, Saliency Maps, Token-level Attribution, Neural Ranking.

1 Introduction

Composed Image Retrieval (CIR) is the task of retrieving a ranked list of images that resemble a reference image but differ in a specific way described by a natural language prompt [43]. Unlike traditional image search, which typically relies on either visual similarity or textual queries alone, CIR fuses both modalities into a single query. For example, a user may ask for images similar to a photo of a forest, but "in autumn colors" or "at night". Solving this task requires the system to interpret the transformation described in the prompt and apply it to the visual concept. CIR has gained attention in multimedia retrieval and computer vision [37, 28], especially with the introduction of vision-language models such as CLIP [30], which embeds both images and text into a joint feature space. Joint image-text embeddings form the basis of many recent CIR models, including zero-shot systems that leverage pre-trained components without further task-specific training [12, 4, 50].

In practice, CIR remains difficult to work with. Existing workflows rely heavily on "blind" trial-and-error: developers diagnose a low Recall@K metric and randomly tweak the prompt or reference image without understanding the underlying failure mechanism. While general-purpose tools exist for embedding visualization [36] and unimodal explainability [34], they are insufficient for CIR for two distinct reasons. First, generic embedding projectors treat queries as static data points, failing to visualize the vector shift induced by the textual modifier, the core mechanism of CIR. Second, standard saliency maps often visualize visual features in isolation, failing to reveal when and why the retrieval model ignores the textual condition in favor of visual dominance (e.g., retrieving a "green apple" despite the prompt "red"). Furthermore, as we observe in our analysis, standard dimensionality reduction techniques like UMAP can suffer from style bias in CLIP-based spaces, where images are clustered by artistic style rather than the semantic categories relevant to the user. These limitations hinder both system debugging and prompt engineering, especially in applications where precise control over image semantics is needed [38].

To address this gap, we present InfoCIR, a visual analytics system designed to support the exploration, diagnosis, and enhancement of CIR models. Developed following a design study methodology [33], InfoCIR targets developers and researchers who face recurrent challenges in diagnosing retrieval failures and iteratively refining multimodal queries.

As illustrated in Fig. 1, the interface integrates six linked components (A-F) that support interactive exploration. Users can select a reference image and a text prompt, view top-k retrieval results, inspect their distribution in the embedding space through Uniform Manifold Approximation and Projection (UMAP) [24], and identify key factors that influence model behavior. The system overlays saliency maps, token attributions, and prompt variants to help users understand and refine the search process. See the caption of Fig. 1 for a detailed breakdown of each component.

We present InfoCIR explicitly as a system contribution, emphasizing completeness and practicality in bridging the gap between black-box retrieval models and human intent. Our goal is to operationalize retrieval, explainability, and prompt engineering into a coherent workflow that supports human-in-the-loop reasoning. Our specific contributions include:

- A unified visual analytics architecture that integrates retrieval, attribution, and prompt engineering, enabling a "diagnose-compare-enhance" feedback loop to address the limitations of isolated, "blind" trial-and-error workflows.
- An interactive relevance-feedback workflow that couples LLM-generated prompt variants with visual verification via saliency maps, token-level attribution, and a novel Ideal-anchored Rank-Δ heatmap to verify prompt impact.
- A supervised metric projection pipeline utilizing a debiased UMAP approach to project results, reference images, and queries into a common space, specifically designed to overcome style bias for more accurate neighborhood analysis.
- A controlled user study demonstrating that InfoCIR significantly improves retrieval efficiency and success rates compared to a baseline CIR interface.

Together, these features turn InfoCIR into a practical tool for diagnosing retrieval failures, understanding model behavior, and refining CIR workflows in an interpretable and interactive manner.

2 Related Work

This section reviews five research areas that shape the design and functionality of our system: (1) **Visual analytics for retrieval systems** provide interactive tools for examining ranking dynamics, analyzing query effects and refining results. (2) **Dimensionality reduction and embedding visualization** project high-dimensional embeddings into navigable spaces, allowing users to inspect structural patterns and model behavior. (3) **Visual saliency for retrieval explanation** highlights image regions and textual cues that most influence similarity scores, revealing the model's decision process. (4) **Prompt enhancement and token attribution** link ranking changes to individual words and offer targeted, interpretable suggestions for query refinement. (5) **Positioning of our work** shows how the system unifies and extends these strands to deliver an explainable and interactive CIR workflow.

2.1 Visual Analytics for Retrieval Systems

Recent work in retrieval has increasingly emphasized interactive visual analytics to reveal ranking behavior, query formulation effects, and feedback loops. Early systems such as RankExplorer [35] visualized how rankings evolve over time, while Colorslope [44] extended this concept with slope graphs and heatmaps to capture dense ranking trajectories. These tools laid the foundation for systems that help users interpret the dynamic behavior of rankers in response to changing queries.

Some systems have explored prompt-based interaction at a higher level. PromptMagician [13] supports keyword recommendation and multilevel prompt editing for text-to-image generation, offering a practical interface for experimenting with compositional queries.

Another important direction is counterfactual analysis. CREDENCE [31] introduces query and document edits to explain ranking shifts, extended by recent frameworks for editable search explanations [49, 5]. In image retrieval, even a single well-chosen keyword has been shown to improve CLIP-based results significantly, without requiring retraining [25].

These developments illustrate a shift away from static evaluation toward interactive systems that help users analyze ranking dynamics, explore prompt effects, and understand retrieval behavior-principles that our system adopts and extends.

2.2 Dimensionality Reduction and Embedding Visualization

Visualizing high-dimensional data in a meaningful way requires dimensionality reduction techniques that preserve the structure of

the original space. Classical methods such as Principal Component Analysis (PCA) [18] and Isometric Mapping (Isomap) [39] reduce dimensionality by preserving variance or geodesic distances, respectively. More recent methods such as t-distributed Stochastic Neighbor Embedding (t-SNE) and UMAP emphasize local neighborhood preservation and scalability, making them suitable for large-scale vision-language embeddings [41, 24]. Other techniques including Triplet Map (TriMap) [3], which uses triplet constraints, and Pairwise Controlled Manifold Approximation and Projection (PaCMAP) [46], which balances global and local structure through staged optimization, have expanded the design space for projection-based analysis. These approaches have been applied in multimedia tasks, ranging from interactive music exploration [40] and fast image annotation [23], to clustering validation in dense text retrieval [22]. In our system, we adopt this strategy by using an adjusted UMAP projection to visualize the embedding space in a single, interpretable view that supports retrieval-driven diagnostics.

Building on these dimensionality reduction methods, embedding visualization tools allow users to explore and compare model representations interactively. TensorBoard's Embedding Projector introduced interactive PCA, t-SNE and UMAP plots with brushing and label coloring, providing one of the first practical interfaces for high-dimensional embedding analysis [36]. EmbeddingVis expanded this functionality to support side-by-side model comparison [21], while WizMap scaled these interactions to millions of points using quadtree summaries and map-style navigation [47]. Despite their utility, most of these tools remain model-agnostic and task-agnostic, treating each embedding as an isolated point. They lack retrieval-specific signals such as rank position, prompt evolution, and ideal-image anchors, all of which are integrated into our system [17, 53].

2.3 Visual Saliency for Retrieval Explanation

Visual saliency helps users interpret why a retrieval model ranks certain images as relevant, by highlighting the most influential regions or features. Similarity-based methods such as [10] perturb candidate images to identify areas that drive similarity scores, while interactive frameworks like [42] integrate these maps into feedback loops that enable real-time prompt or crop adjustments. Gradient-based techniques extend saliency to multimodal retrieval: Grad-ECLIP [54] adapts Grad-CAM [34] to CLIP by back-propagating through both image and text towers to produce joint heatmaps and token-level gradients, and SANE [27] links patch-level saliency to semantic attributes for improved interpretability. In the context of CIR, where queries combine visual and textual cues, such techniques help users understand how models align the modalities. Our system incorporates and adapts Grad-ECLIP to support post-retrieval explanation: after each search, a heatmap is overlaid on the reference or any retrieved image, connecting salient regions to influential prompt tokens and guiding targeted prompt enhancement.

2.4 Prompt Enhancement and Token Attribution

Early counterfactual studies showed that small query edits can reveal why certain items appear at incorrect ranks [5]. PromptBench [55] and Counterfact [7] extend this idea to language models by attributing rank shifts to individual tokens instead of entire prompts.

In the context of multimodal models, gradient-based methods like Grad-ECLIP, SANE, and PromptExp compute signed influence scores for each sub-word across CLIP, image-text, and LLM tasks respectively [54, 27, 11]. These studies show that token-level signals more effectively reveal issues such as hallucinated attributes, compared to prompt-level saliency.

While visual summarization tools like RankExplorer [35] and Colorslope [44] visualize rank shifts under prompt edits, they do not link changes to token-level causes. We adopt their color-encoded Rank-Δ design, but couple every row with a bar chart of per-token gradients, closing the causality loop between what moved and which word caused the move.

By combining token attribution with rank-change visualization, our system extends explanation techniques originally developed for text-only models to the multimodal setting of CIR. This integration bridges counterfactual query editing with saliency-based interpretation, enabling users to trace ranking changes back to specific words and their corresponding visual evidence.

2.5 Positioning of our work

Existing CIR models [1, 12] often focus on isolated components, overlooking the need for an integrated approach. Visual embedding tools typically ignore retrieval context and prompt semantics, while ranking interfaces emphasize accuracy without offering interpretability. Dimensionality reduction methods remain model-agnostic, providing limited insight into retrieval behavior, and prompt enhancement tools usually support only a single modality with minimal explanation.

Our system addresses these gaps by unifying retrieval ranks, prompt evolution, and token attribution in a single analytics platform, enabling direct comparison across CIR models. To our knowledge, it is the first system to jointly support analysis of embedding space, query dynamics, and model behavior.

3 Design Rationale & Goals

The design of InfoCIR is rooted in the need to understand complex model behavior when queries are dynamically changing during the retrieval process. Rather than treating the retrieval as a static operation, our system views it as an iterative dialogue where the user refines multimodal inputs to align the model's internal representation with their specific intent. This section formalizes the retrieval and enhancement tasks and outlines the goals that guide our system's interactive capabilities.

3.1 Formal Task

Composed Image Retrieval (CIR) targets the problem of retrieving images that resemble a given reference image but differ according to a textual modification. Unlike traditional image retrieval, CIR requires the system to interpret a natural language transformation and apply it to the visual input, fusing both modalities into a single, semantically meaningful query. This process often involves mapping the reference image into the natural language embedding space, a process known as textual inversion, to allow for a seamless fusion with the text prompt.

Formally, given:

- a reference image $I_{\text{ref}} \in \mathscr{D}$ (query image),
- a natural language modifier T (relative caption),
- a database $\mathscr{D}$ of images with precomputed embeddings,

the goal is to retrieve a ranked list of the top-k target images that are similar to I_{ref}, except for the modifications described by T. This is operationalized as a nearest-neighbor search in a joint multimodal embedding space:

$$\left(I_{(1)}, \ldots, I_{(k)}\right) = \underset{I_j \in \mathscr{D}}{\operatorname{argsort}}\ \operatorname{sim}\Big(f_c(I_{\text{ref}}, T), f_i(I_j)\Big),$$

where f_c is a composition function that fuses the image and text into a query embedding, and f_i is a fixed encoder for the database images.

Prompt enhancement is the task of rewriting the original text modifier T so that the composed query embeds more accurately the user's intent. After the initial search, the user selects one or more ideal images I^* from the top-k results. A language-model-based editor $\mathscr{R}$ then produces a small set of alternative prompts conditioned on the visual context of the reference image I_{ref}, the original text modifier T, and the semantic content of the chosen ideal image(s):

$$T' = \mathscr{R}\big(I_{\text{ref}}, T, I^*\big).$$

Each candidate prompt T' is re-evaluated by the retrieval engine to determine its impact on the ranking of I^* relative to the original query. This single-pass rewrite aims to boost retrieval quality by better aligning the textual modifier with the visual features of the intended target without requiring model retraining.

3.2 Interaction Analysis

To motivate the design of our visual analytics system for CIR, we analyze how users reason while iteratively composing and refining multimodal queries. Rather than issuing a single query and accepting the result, users engage in a continuous sensemaking loop, modeled after Pike et al.'s science of interaction [26], where they must diagnose retrieval failures, compare potential solutions, and enhance their queries to improve outcomes. Within this loop, each interaction functions as a micro-scale hypothesis test that updates the user's mental model of how the system processes multimodal inputs. This structured "diagnose-compare-enhance" workflow ensures that interaction leads to both better retrieval results and a deeper understanding of model behavior.

Throughout this process, users pursue a set of recurring high-level interaction intents that structure how CIR results are analyzed. These intents describe users' macro-level objectives and reflect the cognitive stages of sensemaking described by Pike et al. We characterize these objectives within the three main elements of our proposed workflow:

- Diagnose: Users examine why the system behaves as observed to identify retrieval failures. This involves assessing the semantic coherence of the results, identifying structural outliers in the distribution, and detecting over-represented concepts that may be biasing the retrieval.
- Compare: Users evaluate different states of the retrieval engine to determine which formulation is most effective. This includes switching between prompt variants to observe which produces more coherent semantic neighborhoods and comparing the positional shifts of target images in the embedding space.
- Enhance: Users iteratively modify the composed prompt to better express the target concept. Each revision represents a new hypothesis that is validated by observing rank changes of target images and leveraging automated suggestions to refine the search focus.

These elements are tightly coupled and typically occur within a single sensemaking episode, with users moving fluidly between diagnosis, comparison, and enhancement as they iteratively improve CIR outcomes.

3.3 Design Goals

Building on the "diagnose-compare-enhance" workflow defined above, we formulate three design goals to operationalize these cognitive stages into specific system capabilities. These goals guide the interface design, ensuring that every visual component directly supports a distinct phase of the user's reasoning process.

DG1 Diagnose Support rapid assessment and diagnosis of retrieval behavior. The system must provide interpretable visual signals regarding both global semantic structure and local model decisions, enabling users to quickly judge retrieval relevance and pinpoint the causes of reasoning breakdowns.

DG2 Compare Enable comparison across prompt variants and retrieval outcomes. The system should explicitly visualize shifts in ranking dynamics and semantic proximity between queries, allowing users to contrast outcomes and determine which formulation best captures the intended concept.

DG3 Enhance Support iterative prompt enhancement. The system must facilitate rapid experimentation by providing intelligent query guidance and immediate feedback on refinements, enabling users to efficiently test new hypotheses without blind trial-and-error.

4 Methodology

The proposed interactive multimedia analytics framework for Composed Image Retrieval (CIR) combines multimodal deep learning with coordinated visual representations to support prompt refinement and model introspection. It centers around user feedback, interpretable visualizations and prompt enhancement strategies, leveraging large language models (LLMs). The system follows a modular pipeline architecture, shown in Fig. 2, inspired by the work of Worring et al. [48], with distinct stages for input, model selection, retrieval, projection, visualization and prompt enhancement. The methodological components of the system are described below.

4.1 System Overview

The pipeline begins with a user-provided input query consisting of a reference image and an accompanying natural language prompt. It integrates the SEARLE [1] CIR model, which uses CLIP to compute joint image-text embeddings for comparative retrieval and visualization workflows.

These embeddings are used to rank images in a target database by similarity. The top-k most similar results are visualized using a suite of coordinated views including a 2D embedding view, image gallery, class distribution histogram, and semantic word cloud. An integrated explainable prompt enhancement module uses an LLM to generate alternative text variants, re-ranks them against a user-chosen ideal image, and instantly visualizes their impact through saliency maps, compact Rank-Δ heatmaps and token-influence summaries, giving users a rapid, transparent loop for refining their queries.

The system is implemented using the Plotly and Dash frameworks, which support highly interactive, web-based visualizations with seamless integration of Python-based machine learning components, as illustrated in the system overview image (Fig. 1).

4.2 Input and Model Selection

The system accepts two primary inputs: (1) a reference image which is uploaded and (2) a textual prompt describing the desired transformation or characteristics. A third, optional (but strongly recommended) input is the ideal image(s): from the initial top-k list the user selects one or more images that align with their intent. These images, become the fixed anchor, against which all prompt variants are evaluated, enabling fine-grained counterfactual diagnostics described in Sec. 4.6.

For the retrieval back-end, InfoCIR implements SEARLE, a zero-shot model that performs textual inversion by transforming the reference image into a pseudo-token within CLIP's text embedding space. This architecture enables the compositional fusion of the image and prompt into a joint query vector. The system then calculates cosine similarities between this vector and pre-computed database embeddings to generate the final image rankings.

4.3 Retrieval and Similarity Computation

Given the composed query embedding, the system ranks all database images using cosine similarity. The top-k images are selected for further visualization, forming a ranked result set. Alongside the images, similarity scores and, when available, associated class labels or tags are returned for further analysis.

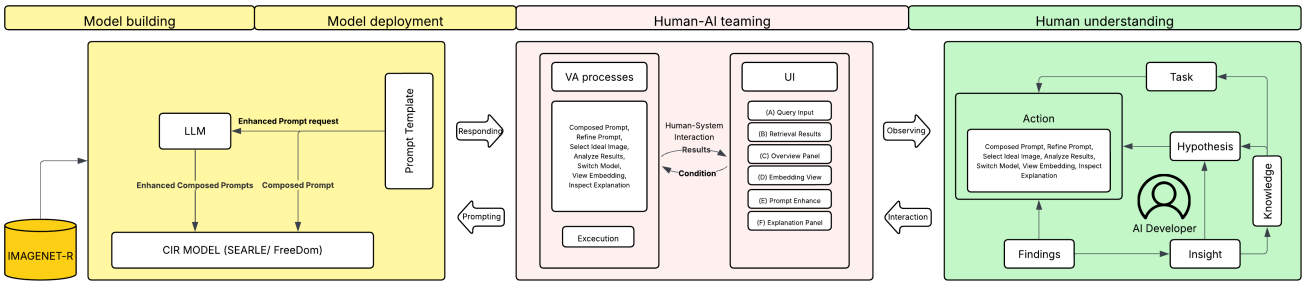

Figure 2: The architecture links prompt composition, embedding-based retrieval, and interactive visualization across UI panels (A-F), enabling users to enhance prompts, interpret outputs, and derive insights. Following Wang et al. [45], we integrate semantic fusion (CLIP space), visual integration (saliency and token overlays), and cross-view linking (synchronized views).

4.4 Dimensionality Reduction

To facilitate qualitative, topological inspection of the embedding structure, the high-dimensional vectors of the top-k results are projected to a two-dimensional space. This projection is intended to support diagnosis through visual neighborhood validation, rather than precise metric reasoning over embedding distances.

In this domain, UMAP and t-SNE stand out as the standard techniques for high-dimensional visualization. We employ UMAP due to three workflow-driven requirements:

1. **Parametric Stability (Out-of-Sample Extension):** Our system embeds new prompts into a fixed reference map in real time. UMAP's parametric formulation supports stable transform operations on unseen data [32], whereas standard t-SNE is non-parametric and typically requires re-running the projection.

2. **Global structure preservation:** Like t-SNE, UMAP emphasizes local neighborhood preservation, which is essential for identifying semantically related groups of embeddings [24]. However, UMAP typically produces more coherent coarse-grained groups of clusters [8].

3. **Cluster Density for Rapid Semantic Verification:** UMAP tends to produce denser and more compact clusters with clearer inter-class separation [19]. This directly supports our DG1 of enabling users to rapidly verify whether a prompt has landed within the intended semantic category.

The resulting two-dimensional coordinates are used to position images in the Embedding View.

To ensure the visualization supports semantic debugging, we apply a Class Projection Preservation Pipeline (Fig. 3) prior to the final projection. Raw CLIP embeddings often exhibit significant overlap between classes due to stylistic or background features. We explicitly frame this pipeline not as model training or adaptation, but as a heuristic, task-driven semantic pre-alignment strategy applied before visualization.

The pipeline consists of three sequential filtering steps: (A) **Style Debiasing** via PCA-based filtering to retain dimensions with higher semantic discriminability; (B) **Contrastive Debiasing** [15], which encourages embeddings to align more closely with their class prototypes; and (C) **Independent Component Analysis (ICA)** [29] to separate semantic signals from stylistic noise.

This heuristic pipeline aims to improve class-level separability so that the resulting visualization reflects user-relevant semantic categories rather than irrelevant stylistic variation. Finally, we underline that the Panel D is intended for qualitative, topological inspection rather than precise metric reasoning. Distances and directions in the two-dimensional projection should not be interpreted as faithful representations of semantic magnitude or linear embedding trajectories.

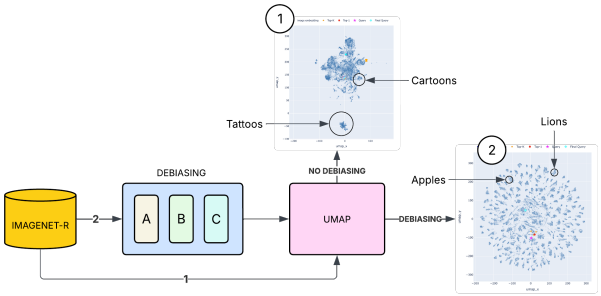

Figure 3: Conceptual overview of the Class Projection Preservation Pipeline. ImageNet-R embeddings are passed through a three-stage filtering process: (A) Style Debiasing, (B) Contrastive Debiasing, and (C) Alternative Projection via ICA, prior to the final UMAP visualization. Compared to the unprocessed projection (1), the processed version (2) exhibits improved class separability and reduced influence of style artifacts, aiding visual verification of class membership.

4.5 Interactive Visualization

The retrieved and projected images are rendered in six main panels (A–F) as illustrated in Fig. 1, providing multiple coordinated views to support the diagnose-compare-enhance workflow and fulfill the design goals:

- **Composed Image Retrieval Panel (A):** Allows users to input a reference image and a text prompt while selecting the number of images k to be retrieved.

- **Query Results Panel (B):** Displays the retrieved images in an Image Gallery sorted by similarity to the query. This view supports rapid inspection of ranked outputs (DG1) and is synchronized with the Embedding View.

- **Histogram / Word Cloud Panel (C):** Includes a class-frequency Histogram (C_1) to analyze category dominance and a Word Cloud (C_2) that summarizes semantic trends within the current top-k. These widgets help users diagnose over-represented concepts or bias in the results, explicitly addressing DG1.

- **Embedding View (D):** The central panel shows a 2D UMAP projection of the dataset. Spatial proximity reflects coarse se-

mantic neighborhood relationships rather than precise similarity magnitudes, allowing users to observe class-level clustering, separation, or outliers (DG1). Interactive selection reveals metadata associated with the respective projection.

- **Prompt Enhancement Panel (E):** Proposes alternative prompt variants T' generated by an LLM, conditioned on user-selected ideal images. Users can compare these variants interactively to observe their impact on retrieval performance, facilitating comparison (DG2) and iterative refinement (DG3).
- **Explanation Panel (F):** Provides multi-level attribution to expose the causes of model behavior, supporting diagnostic reasoning (DG1). This includes:
  - **Saliency Map (F1):** Overlays gradient-based visual saliency to highlight influential image regions, linking visual evidence to textual input and assisting users in understanding model focus.
  - **Token Attribution Chart (F2):** Quantifies how strongly each token in the prompt contributes to the retrieval of the ideal image.
  - **Rank-Δ Heatmap (F3):** Visualizes how each variant affects the position of the originally retrieved top-k images. Cell colors show rank changes (green for upward moves, red for drops), allowing users to enhance the search (DG3) by identifying the most effective prompts through direct comparison (DG2).

These panels are linked to support multi-faceted exploration. Any interaction in one view, such as clicking an image in the Query Results Panel, automatically updates the associated analytical and explanatory views.

4.6 Explainable Prompt Enhancement

The prompt enhancement module in InfoCIR is a user-guided implementation of the refinement function $\mathscr{R}$. A specific challenge in this workflow is the "cold start" problem, where the initial query fails to retrieve any relevant images in the top-k to serve as feedback anchors. InfoCIR addresses this by utilizing the manual Composed Image Retrieval Panel (Panel A) as a primary pivoting tool: in zero-hit scenarios, users first adjust the initial modifier T to shift the search into a relevant semantic neighborhood. Once a viable target appears in the top-k, the user selects one or more ideal target images I^*. These chosen images serve as the anchors for enhancement: the system leverages the open-weight mistralai/Mistral-7B-Instruct-v0.2 large language model (LLM) to generate a set of alternative prompts T' grounded on the content of I_{ref}, T, and the visual features of I^*. For each proposed variant T', the system re-ranks the results and visualizes the rank change of I^* via an Ideal-anchored Rank-Δ heatmap. This heatmap provides immediate feedback on how each new prompt affects the target's position, helping users compare prompt efficacy at a glance. Because the heatmap fixes the initial top-k as a stable reference set, it visualizes local re-ordering to preserve visual continuity across variants.

Along with the prompt suggestions, the module provides tools to diagnose why certain prompts perform better. A token attribution chart (using gradient-based analysis [11]) highlights which words in T' have the strongest impact on the similarity score $s = \text{sim}(f_c(I_{\text{ref}}, T'), f_i(I^*))$. Technically, we compute the gradient of s with respect to each input token. The magnitude of these gradients indicates each word's contribution to the match, which is visualized as a bar chart to guide users toward high-influence terms. This offers insight into the model's reasoning about the text query and suggests which keywords might be emphasized or adjusted in further refinements (see Fig. 4).

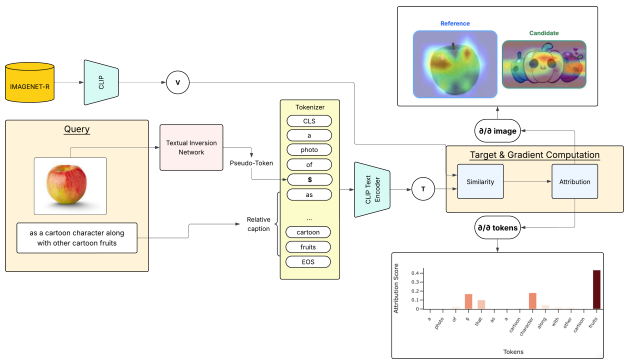

Figure 4: Explanation pipeline for visual and textual attribution in CIR. The system accepts a reference image I_r and a relative caption C, combining them into the prompt "a photo of $\varphi(I_r)$ that C", where $\varphi(I_r)$ is a pseudo-token generated via a textual inversion network. This prompt is tokenized and encoded by CLIP's text encoder to produce a text embedding T. Candidate images from the dataset (e.g., ImageNet-R) are encoded into image embeddings V, and retrieval is based on cosine similarity $s = \cos(V, T)$. To generate explanations, the similarity score s is differentiated with respect to either the image or the text tokens. The resulting gradients yield visual saliency maps (top right) and token attribution scores (bottom right), helping users understand which parts of the image or prompt influenced the retrieval result most.

Complementing the token-level explanation, a saliency map is generated to show which regions of I_{ref} or I^* are most influential. We employ gradient-based visual explanation techniques, and more specifically implement from scratch an adaptation of Grad-ECLIP [54] for our CLIP-based multimodal encoder (SEARLE), to attribute the output of f_c back to image pixels. By back-propagating the gradient of s through the model's image tower, we produce a heatmap overlay denoting areas crucial to the query-result similarity. The saliency highlights enable users to see which parts of the image the model "looked at" when matching the query, linking concrete visual evidence to the semantic query feedback (see Fig. 4).

Together, these explainable prompt enhancement features put the user in control of an interactive feedback loop. The LLM-generated variants T' (with their rank-change indicators) and the attribution visualizations (token bars and saliency maps) collectively, allow users to steer the search with greater transparency and confidence. This approach aligns with emerging interactive explanation paradigms that emphasize human-in-the-loop refinement and deeper model transparency [31], ensuring that prompt adjustments are both informed by the model's inner workings and grounded on the user's intent, ultimately fulfilling DG3 by lowering the barrier to effective experimentation.

4.7 System Modularity and Extensibility

The system is designed with modularity in mind. Retrieval models can be easily swapped or added as long as they provide image and text embeddings in the CLIP embedding space. Visualization modules are decoupled and configurable, supporting the addition of new views or encoding strategies. This flexibility allows for experimentation with retrieval algorithms, datasets and user interaction paradigms.

4.8 Usage Scenario

To demonstrate how InfoCIR supports interactive refinement in CIR, consider a user uploading an image of a green apple and entering the prompt "a red apple". The goal is to retrieve semantically

relevant images that preserve the object identity (apple) while modifying its attribute (color).

Upon querying, the SEARLE retriever returns a result set that includes both green and red apples, along with unrelated fruits. In the UMAP projection, results are scattered, with no clear cluster for red apples. The user opens the attribution panel, which shows that the term "red" contributes weakly to the ranking, suggesting a style bias in the CLIP representation.

To correct this, the user selects one or more retrieved images that match their intent (e.g., any clearly deep-red apple) and opens the Prompt Enhancement module. This component suggests semantically guided prompt variants, such as "ripe apple", "crimson apple", and "dark red fruit", each grounded on corpus embeddings and visual similarity.

After choosing "crimson apple", the updated retrieval results appear in a compact cluster of red apples in the embedding space. The selected image jumps from rank 12 to rank 2, as shown in the Rank-Δ heatmap, and the attribution weights now reflect stronger activation on "red" and relevant image regions.

This scenario illustrates how InfoCIR enables users to debug and steer the retrieval process through prompt adaptation and embedding inspection, leading to semantically faithful results that align with the composed intent.

5 Evaluation

Our evaluation strategy focuses on the Prompt Enhancement workflow as the primary unit of analysis. While InfoCIR comprises six coordinated views (A-F), prompt enhancement acts as the "keystone" task that necessitates the synthesis of information from the entire dashboard. To successfully refine a query, a user must first diagnose the initial failure (using the Embedding View and Saliency Maps), formulate a hypothesis (using the LLM suggestions), and verify the outcome (using the Rank-Δ heatmap). Consequently, we argue that measuring performance on prompt enhancement serves as a holistic proxy for the utility of the supporting diagnostic components.

We conducted a two-part evaluation to assess how InfoCIR supports users in efficiently retrieving target images and understanding model behavior through its interactive analytical widgets. The study combined quantitative performance tasks with qualitative feedback analysis to examine both objective efficiency and subjective user experience. Our main hypothesis was that InfoCIR's interactive explanatory features, particularly the Prompt Enhancement Panel, would improve retrieval performance, user confidence, and understanding of image-text relationships compared to a baseline system without such support.

This study adhered to established ethical research standards. All participants were informed about the study and provided written consent. Participation was voluntary, no identifying data were collected, and all responses were anonymized and stored securely. In line with GDPR regulations, the data will be permanently deleted after the project is completed.

5.1 Quantitative Study: Performance and Retrieval Efficiency

5.1.1 Procedure

We recruited eight Master's students in Artificial Intelligence from the University of Amsterdam, all with a background in computer science. Participants had previously worked with web development, design, or information retrieval tools, and had experience searching for and retrieving images as part of past professional tasks. Henceforth, we will refer to the participants as P1 through P8.

For the dataset, we used ImageNet-R to simulate realistic retrieval scenarios with similar difficulty and visual abstraction. The reference images for Task A and Task B were obtained from publicly available online sources[1] [2]. Two image retrieval tasks (Task A and Task B) were designed to have comparable semantic and perceptual difficulty. In both tasks, participants viewed a realistic reference image and attempted to retrieve a stylized (cartoon-like) version of the same object category. Task A involved a **porcupine** and Task B a **Boston Terrier** (see Fig. 5). Both categories were selected for their similar visual complexity and distinct, recognizable features.

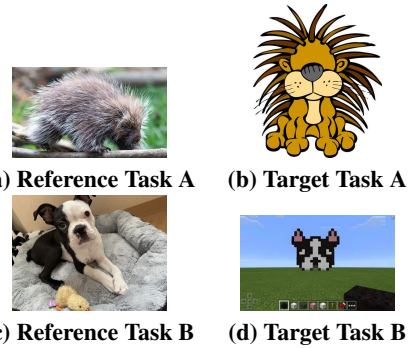

(a) Reference Task A **(b) Target Task A**

(c) Reference Task B **(d) Target Task B**

Figure 5: Reference and target image pairs used in the quantitative evaluation. Task A (porcupine) and Task B (Boston Terrier) each required retrieving a stylized version of a realistic reference image. Participants iteratively refined a text prompt until the target appeared within the top-3 retrieved results.

Design and Procedure We compared participants' performance while using two different versions of the tool, one basic and one enhanced, to see how the widgets affect the image retrieval task.

1. **Scenario 1 - Baseline:** Participants used only the Composed Image Retrieval and Query Results panels, with all InfoCIR's widgets hidden.
2. **Scenario 2 - InfoCIR:** Participants used the full InfoCIR interface, including the Prompt Enhancement Panel, Saliency Map, Histogram/Word Cloud, Embedding View, Token Attribution, and Rank-Δ heatmap.

To avoid order bias, participants were split into two groups that completed the same tasks in reversed order, ensuring task sequence did not influence performance.

- **Group A:** Task A (Baseline) → Task B (InfoCIR)
- **Group B:** Task B (Baseline) → Task A (InfoCIR)

Each task required participants to refine a text prompt until the target appeared in the top-3 retrieved images, with a 5-minute limit. Each query returned ten images. Sessions were timed, and two experimenters recorded results. Performance was evaluated using:

1. Success Rate: proportion of participants who placed the target in the top-3.
2. Average Time on Task: mean time to successful retrieval.
3. Number of Unique Queries: average number of prompt submissions.

5.1.2 Quantitative Findings

As summarized in Tab. 1, participants using the full InfoCIR interface showed markedly improved retrieval performance compared to the baseline. The success rate increased from 37.50% to 87.50%,

[1] Reference image for Task A, retrieved from: https://brandywinezoo.org/wp-content/uploads/2024/10/north-american-porcupine-brandywine-zoo-2023.jpg

[2] Reference image for Task B, retrieved from: https://dogtime.com/wp-content/uploads/sites/12/2023/09/GettyImages-1500664833.jpg

average completion time decreased from 277 seconds to 133 seconds, and the average number of unique queries dropped from 7.4 to 3.3. One participant (P6) did not succeed in either condition, as the initial provided prompts were overly generic and the target did not appear within the top ten retrieved images, leaving no meaningful feedback to guide refinement and highlighting that effective use of InfoCIR still requires purposeful prompt adjustment.

These results indicate that InfoCIR's widgets accelerated convergence by guiding users toward more effective, category-based language. The metrics confirm the fulfillment of the design goals: the reduced time and query counts demonstrate support for rapid experimentation (DG3), while the high success rate shows that users successfully diagnosed and corrected retrieval failures (DG1).

5.2 Qualitative Study: User Experience and Interpretability

5.2.1 Procedure

After completing both retrieval tasks, participants filled out a post-study questionnaire containing five-point Likert scale evaluations of each InfoCIR component. A short discussion followed in which participants described their workflow, what they found intuitive, and where they encountered difficulties. The goal of this qualitative phase was to understand how users experienced the interface while performing retrieval, which components they relied on most, and how the system supported their reasoning.

5.2.2 Qualitative Findings

Participants generally agreed that the enhanced interface improved both retrieval efficiency and their understanding of how the system processes prompts. However, several participants (P1, P3, P5, P8) noted an initial period of orientation, during which they explored the various components before converging on a smaller set of features that they relied on consistently. This indicates a learning curve, especially for components intended to aid interpretability rather than immediate task execution.

The **Prompt Enhancement Panel** was consistently identified as the most useful feature for improving retrieval. The majority of participants rated its suggestions as clear and relevant, and 62.50% of participants selected it as the most helpful widget overall. Participants claimed that "It made the task easier" (P4) and "I do not need to manually craft a strong initial prompt in order to retrieve the target image within the top-3 results" (P7). Moreover, participants indicated that using the panel taught them how to phrase category-based prompts more effectively rather than relying on descriptive attributes of the target image (figure available in supplementary material), validating its role in supporting iterative prompt enhancement (DG3).

The **Histogram and Word Cloud** helped participants identify more effective vocabulary by showing which concepts the model associated with their current prompt. Participants reported that it supported the shift from descriptive phrasing toward category based terms. For example, P2 explained that it helped change "small black and white dog" to "Boston Terrier", which directly improved retrieval, and P5 described it as a useful guide for discovering the model's preferred terminology (figure available in supplementary material). By exposing these semantic mismatches, these widgets successfully fulfilled the goal of supporting rapid diagnosis (DG1). However, participants also noted that while the widget aided understanding, it contributed more to overall orientation than to fast prompt refinement during the task.

Both the **Embedding View** and the **Saliency Map** contributed meaningfully to model interpretability and post hoc reasoning rather than immediate retrieval decisions. Participants generally understood the Embedding View and used it primarily to observe how the system groups images and organizes similarity relationships. Several participants reported that it helped them grasp the model's internal structure but not in a way that directly supported fast prompt refinement. P5 noted that it "explains the structure, but I cannot use it quickly". Concurrently, the Saliency Map assisted users by verifying which image regions influenced retrieval and was often employed as a correctness check (figure available in supplementary material). P5 described it as useful for confirming whether the system was focusing on the main object, aligning with the diagnostic requirements of DG1. However, participants commonly noted a temporal cost, commenting that it required more time and attention than the task allowed. P3 stated that it "takes too much time to deal with", and P4 mentioned that although it clarified model attention, it did not directly help with prompt adjustment.

The **Token Attribution** and **Rank-Δ Heatmap** were used less frequently during the time-constrained retrieval tasks, as many participants found them slower to interpret and therefore difficult to apply to fast prompt adjustments. However, these widgets supported deeper analysis for participants who adopted a more exploratory approach. P3 and P8, in particular, spent additional time examining these visualizations during Scenario B and reported that they helped them understand how subtle phrasing changes influenced ranking outcomes (figure available in supplementary material). This suggests that these widgets are valuable for detailed interpretability and comparison (DG2) when users have sufficient time to explore them.

6 Discussion and Future Work

The development of InfoCIR highlights the growing need for explainable, interactive systems in the domain of Composed Image Retrieval (CIR). While recent CIR models show promising zero-shot and training-free retrieval capabilities [1, 12], their behavior often remains opaque to users, particularly when multimodal queries (image + text) yield unexpected results. Our system addresses this challenge by providing a tightly integrated workflow, through which users can trace retrieval dynamics through interpretable and visual means.

A key design insight was the benefit of combining retrieval results with both the overall structure of the embedding space and detailed feedback on individual prompt tokens. Embedding projections via UMAP support exploration of neighborhood coherence and cluster structure, while token attribution and saliency maps provide insights on the internal reasoning behind model decisions. The connected and interactive visual views allow users to detect and act on misalignments, such as when a prompt term appears dominant despite being semantically not relevant, making it easier to quickly form and test new ideas.

Another important contribution is the idea of ideal-guided prompt enhancement. By allowing users to fix one or more target images as benchmarks, InfoCIR turns prompt enhancement into a constrained optimization task visible through the Rank-Δ heatmap. This approach offers immediate, localized feedback on prompt effectiveness, addressing a longstanding gap in model debugging workflows, where prompt quality is typically judged only by final retrieval rank without insight into causal dynamics [6, 20].

However, limitations remain:

- **Limitations of Embedding Visualization** Projecting high-dimensional CLIP embeddings into 2D inherently introduces metric distortion, meaning that linear vector arithmetic valid in the original space does not strictly hold in the projection. Consequently, the Embedding View is intended for visual cluster verification, confirming topological containment within a semantic cluster, rather than as a precise metric map for trajectory analysis. While our pipeline facilitates this task by creating distinct clusters for rapid visual inspection, we acknowledge the inherent trade-offs in manifold approximation. Future work will explore alternative techniques such as PaCMAP and TriMap [46, 3, 14] to potentially achieve a better balance between local fidelity and global structural con-

Table 1: Participant performance across the two retrieval scenarios. Scenario A is the baseline (manual prompt refinement). Scenario B is the InfoCIR-assisted condition (all analytical widgets and the Prompt Enhancement Panel enabled). Success indicates whether the target image appeared in the top-3 retrieval results. For Scenario B, each prompt refinement action was also counted as a query. Averages are across all 8 participants.

		Scenario A (Baseline)			Scenario B (InfoCIR-Assisted)		
ID	**Group**	**Time (s)**	**Queries**	**Success Rate (%)**	**Time (s)**	**Queries**	**Success Rate (%)**
P1	A	243	6	100.00	103	3	100.00
P2	A	300	9	0.00	87	2	100.00
P3	A	300	7	0.00	134	3	100.00
P4	A	281	5	100.00	91	3	100.00
P5	B	300	11	0.00	149	4	100.00
P6	B	300	8	0.00	300	6	0.00
P7	B	194	4	100.00	55	2	100.00
P8	B	300	9	0.00	142	3	100.00
Average	–	**277**	**7.4**	**37.50**	**133**	**3.3**	**87.50**

straints.

- **Cold Start Vulnerability:** The automated prompt enhancement relies on the user selecting at least one relevant "ideal" image from the initial results. This creates a dependency on the initial query performance. If the zero-shot retrieval fails to return any relevant targets (a "zero-hit" scenario), the user must rely on manual iteration to pivot the search into a relevant semantic neighborhood before the automated enhancement tools can be engaged.
- **Generalization vs. Overfitting:** By anchoring prompt refinement to specific user-selected images, there is a risk that the generated prompts may overfit to the visual features of the chosen anchors rather than generalizing to the broader semantic category. A prompt that successfully retrieves the specific anchor might inadvertently exclude other relevant variations of the target class.
- **Model Actionability:** While InfoCIR effectively diagnoses retrieval failures, the insights are primarily actionable for prompt engineering and query weighting rather than model correction. In the era of large foundation models, developers rarely have the capacity to retrain the heavy backbone (e.g., SEARLE or CLIP) based on granular, instance-level observations exposed by the tool.
- **LLM Heuristics and Verification:** Prompt suggestions are generated by an open-weight LLM and remain heuristic; they do not guarantee retrieval improvement and may occasionally hallucinate attributes. We mitigate this via the Rank-Δ heatmap, which acts as a verification loop, allowing users to empirically test and discard ineffective variants. Future work could incorporate structured prompts (e.g., scene graphs [52]) or multi-turn refinement [51] to improve generation stability [9].
- **Attribution Fidelity:** Token attribution charts depend on gradient approximations that may not fully capture the complex attention mechanisms in transformer-based architectures [16].
- **Local vs. Global Explainability:** Attribution-based explanations offer only local insight into specific query-image pairs. Extending the system to support counterfactual visual generation (e.g., what-if edits to the query image) or concept-level abstractions could provide global interpretability and support model debugging at a higher level [2].
- **Feedback Loop Integration:** Current refinement consists of a single LLM-based rewrite guided only by the user-selected ideal images. Saliency maps and token-attribution bars serve purely as explanatory overlays and are not fed back into the rewrite function. Future iterations could let explicit user annotations, such as region sketches, custom attention masks, or reinforcement-guided scoring, directly influence the prompt generator, providing finer control and richer engagement.

Despite these constraints, InfoCIR demonstrates that aligning retrieval outcomes with human-centered reasoning is both feasible and impactful. The system advances CIR tooling beyond black-box metrics toward interactive, explanation-rich workflows that better support model development, evaluation and refinement.

7 Conclusion

We presented InfoCIR, a visual analytics system designed to enhance transparency, interpretability and refinement in Composed Image Retrieval (CIR). By integrating retrieval ranking, embedding projection and visualizations, prompt enhancement and multi-level explanation within a coordinated and interactive interface, our system addresses key gaps in CIR interpretability and supports developers in exploring, understanding, improving and validating retrieval model behavior with greater confidence.

Our contributions include an explainable prompt-enhancement pipeline that combines LLM-generated prompt variants with gradient-based attribution visualizations, a coordinated environment for inspecting embedding dynamics, and a modular architecture that can accommodate new CIR models with minimal effort. Through guided interface interactions, users can rapidly assess prompt impact, explore the embedding landscape, and pinpoint failure cases or performance bottlenecks.

In contrast to existing tools that isolate embedding inspection or prompt engineering, InfoCIR combines both, supporting a new paradigm of retrieval-centered multimodal analysis.

Future work can extend model support, scalability and personalization, further bridging the gap between opaque model internals and user-centered retrieval workflows. As the field of multimodal AI continues to evolve, we see InfoCIR as both a diagnostic toolkit and a framework for future systems that prioritize human-guided refinement, transparency and trust in retrieval pipelines.

Supplemental Materials

Supplemental materials have been submitted and include (1) An appendix with the UMAP-adjusted Results and the Widgets during the Evaluation Process, (2) a demo of the InfoCIR system, and (3) a hyperlink to the GitHub repository containing the source code of the InfoCIR system.

References

[1] L. Agnolucci, A. Baldrati, M. Bertini, and A. D. Bimbo. isearle: Improving textual inversion for zero-shot composed image retrieval, 2024. 1, 3, 4, 8

[2] A. L. Alfeo, A. G. Zippo, V. Catrambone, M. G. Cimino, N. Toschi, and G. Valenza. From local counterfactuals to global feature importance: efficient, robust, and model-agnostic explanations for brain connectivity networks. *Computer Methods and Programs in Biomedicine*, 236:107550, 2023. doi: 10.1016/j.cmpb.2023.107550 9

[3] E. Amid and M. K. Warmuth. Trimap: Large-scale dimensionality reduction using triplets, 2022. 3, 8

[4] A. Baldrati, L. Agnolucci, M. Bertini, and A. D. Bimbo. Zero-shot composed image retrieval with textual inversion, 2023. 2

[5] B. Chandna and P. Sen. A counterfactual explanation framework for retrieval models. *CoRR*, abs/2409.00860, 2024. 2, 3

[6] B. Chandna and P. Sen. A counterfactual explanation framework for retrieval models, 2025. 8

[7] Z. Chen, F. Silvestri, J. Wang, Y. Zhang, Z. Huang, H. Ahn, and G. Tolomei. Grease: Generate factual and counterfactual explanations for gnn-based recommendations, 2022. 3

[8] Y. chin Ivan Chang. A survey: Potential dimensionality reduction methods, 2025. 5

[9] J. Choi. Efficient prompt optimization for relevance evaluation via llm-based confusion matrix feedback. *Applied Sciences*, 15(9), 2025. doi: 10.3390/app15095198 9

[10] B. Dong, R. Collins, and A. Hoogs. Explainability for content-based image retrieval. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*, pp. 95–98, 2019. 3

[11] X. Dong, S. Wang, D. Lin, G. K. Rajbahadur, B. Zhou, S. Liu, and A. E. Hassan. Promptexp: Multi granularity prompt explanation of large language models. *CoRR*, abs/2410.13073, 2024. 3, 6

[12] N. Efthymiadis, B. Psomas, Z. Laskar, K. Karantzalos, Y. Avrithis, O. Chum, and G. Tolias. Composed image retrieval for training-free domain conversion, 2024. 2, 3, 8

[13] Y. Feng, X. Wang, K. K. Wong, S. Wang, Y. Lu, M. Zhu, B. Wang, and W. Chen. Promptmagician: Interactive prompt engineering for text to image creation. *IEEE Transactions on Visualization and Computer Graphics*, 30(1):295–305, 2024. doi: 10.1109/TVCG.2023.3327168 2

[14] J. Fu and D. Zhao. Mpad: A new dimension-reduction method for preserving nearest neighbors in high-dimensional vector search, 2025. 8

[15] T. Ghosh and M. Kirby. Supervised dimensionality reduction and visualization using centroid-encoder, 2020. 5

[16] S. M. V. Hatefi, M. Dreyer, R. Achtibat, T. Wiegand, W. Samek, and S. Lapuschkin. Pruning by explaining revisited: Optimizing attribution methods to prune cnns and transformers, 2024. 9

[17] Z. Huang, D. Witschard, K. Kucher, and A. Kerren. Va + embeddings star: A state-of-the-art report on the use of embeddings in visual analytics. *Computer Graphics Forum*, 42:539–571, 06 2023. doi: 10.1111/cgf.14859 3

[18] I. T. Jolliffe. *Principal Component Analysis*. Springer Series in Statistics. Springer-Verlag, New York, 2002. doi: 10.1007/b98835 3

[19] D. Kobak and G. Linderman. Initialization is critical for preserving global data structure in both t-sne and umap. *Nature Biotechnology*, 39:1–2, 02 2021. doi: 10.1038/s41587-020-00809-z 5

[20] Q. R. Lauro, S. Shankar, S. Zeighami, and A. Parameswaran. Rag without the lag: Interactive debugging for retrieval-augmented generation pipelines, 2025. 8

[21] Q. Li, K. S. Njotoprawiro, H. Haleem, Q. Chen, C. Yi, and X. Ma. Embeddingvis: A visual analytics approach to comparative network embedding inspection, 2018. 3

[22] Z. Liu, H. Zhang, C. Xiong, Z. Liu, Y. Gu, and X. Li. Dimension reduction for efficient dense retrieval via conditional autoencoder, 2022. 3

[23] F. Luus, N. Khan, and I. Akhalwaya. Interactive supervision with t-sne. In *Proceedings of the 10th International Conference on Knowledge Capture*, K-CAP '19, p. 85–92. Association for Computing Machinery, New York, NY, USA, 2019. doi: 10.1145/3360901.3364414 3

[24] L. McInnes, J. Healy, and J. Melville. Umap: Uniform manifold approximation and projection for dimension reduction, 2020. 2, 3, 5

[25] R. Nara, Y. C. Lin, Y. Nozawa, Y. Ng, G. Itoh, O. Torii, and Y. Matsui. Revisiting relevance feedback for clip based interactive image retrieval. *CoRR*, abs/2404.16398, 2024. 2

[26] W. A. Pike, J. Stasko, R. Chang, and T. A. O'Connell. The science of interaction. *Information Visualization*, 8(4):263–274, Dec. 2009. doi: 10.1057/ivs.2009.22 4

[27] B. A. Plummer, M. I. Vasileva, V. Petsiuk, K. Saenko, and D. Forsyth. Why do these match? explaining the behavior of image similarity models. *arXiv*, 2019. 3

[28] B. Psomas, I. Kakogeorgiou, N. Efthymiadis, G. Tolias, O. Chum, Y. Avrithis, and K. Karantzalos. Composed image retrieval for remote sensing, 2024. 2

[29] B. Póczos. Independent component analysis. https://www.cs.cmu.edu/~bapoczos/Classes/ML10715_2015Fall/slides/ICA.pdf, 2015. Lecture slides, CMU 10-715: Advanced Introduction to Machine Learning. 5

[30] A. Radford, J. W. Kim, C. Hallacy, A. Ramesh, G. Goh, S. Agarwal, G. Sastry, A. Askell, P. Mishkin, J. Clark, G. Krueger, and I. Sutskever. Learning transferable visual models from natural language supervision, 2021. 2

[31] J. Rorseth, P. Godfrey, L. Golab, M. Kargar, D. Srivastava, and J. Szlichta. Credence: Counterfactual explanations for document ranking. In *Proceedings of the 39th International Conference on Data Engineering*, 2023. 2, 6

[32] T. Sainburg, L. McInnes, and T. Gentner. Parametric umap embeddings for representation and semisupervised learning. *Neural Computation*, 33:2881–2907, 10 2021. doi: 10.1162/neco_a_01434 5

[33] M. Sedlmair, M. Meyer, and T. Munzner. Design study methodology: Reflections from the trenches and the stacks. *IEEE Transactions on Visualization and Computer Graphics*, 18(12):2431–2440, 2012. doi: 10.1109/TVCG.2012.213 2

[34] R. R. Selvaraju, M. Cogswell, A. Das, R. Vedantam, D. Parikh, and D. Batra. Grad-cam: Visual explanations from deep networks via gradient-based localization. *International Journal of Computer Vision*, 128(2):336–359, Oct. 2019. doi: 10.1007/s11263-019-01228-7 2, 3

[35] C. Shi, W. Cui, S. Liu, P. Xu, W. Chen, and H. Qu. Rankexplorer: Visualization of ranking changes in large time series data. *IEEE Transactions on Visualization and Computer Graphics*, 18(12):2669–2678, 2012. doi: 10.1109/TVCG.2012.253 2, 3

[36] D. Smilkov, N. Thorat, C. Nicholson, E. Reif, F. B. Viégas, and M. Wattenberg. Embedding projector: Interactive visualization and interpretation of embeddings, 2016. 2, 3

[37] X. Song, H. Lin, H. Wen, B. Hou, M. Xu, and L. Nie. A comprehensive survey on composed image retrieval, 2025. 2

[38] Y. Tang, J. Yu, K. Gai, J. Zhuang, G. Xiong, G. Gou, and Q. Wu. Missing target-relevant information prediction with world model for accurate zero-shot composed image retrieval, 2025. 2

[39] J. B. Tenenbaum, V. De Silva, and J. C. Langford. A global geometric framework for nonlinear dimensionality reduction. *Science*, 290(5500):2319–2323, 2000. doi: 10.1126/science.290.5500.2319 3

[40] P. Tovstogan, X. Serra, and D. Bogdanov. Visualization of deep audio embeddings for music exploration and rediscovery. In *Proceedings of the 19th Sound and Music Computing Conference (SMC)*, pp. 493–500. Saint-Étienne, France, June 2022. Available from UPF Digital Repository. 3

[41] L. van der Maaten and G. Hinton. Visualizing data using t-sne. *Journal of Machine Learning Research*, 9(86):2579–2605, 2008. 3

[42] B. Vasu, B. Hu, B. Dong, R. Collins, and A. Hoogs. Explainable interactive content-based image retrieval. *Applied AI Letters*, 2(4):e41, 2021. doi: 10.1002/ail2.41 3

[43] N. Vo, L. Jiang, C. Sun, K. Murphy, L.-J. Li, L. Fei-Fei, and J. Hays. Composing text and image for image retrieval - an empirical odyssey, 2018. 2

[44] H. Wang, X. Jiang, A. Nagarajan, X. Guo, L. Ding, D. Wan, J. Zhao, and Y. Chen. Colorslope: A balanced visualization of overview and

details on ranks over time. *Visual Intelligence*, 1(7):1–18, 2023. doi: 10.1007/s44267-023-00008-9 2, 3

[45] J. Wang, Z. Deng, D. Deng, X. Wang, R. Sheng, Y. Cai, and H. Qu. Empowering multimodal analysis with visualization: A survey. *Computer Science Review*, 57:100748, 2025. doi: 10.1016/j.cosrev.2025.100748 5

[46] Y. Wang, H. Huang, C. Rudin, and Y. Shaposhnik. Understanding how dimension reduction tools work: An empirical approach to deciphering t-sne, umap, trimap, and pacmap for data visualization, 2021. 3, 8

[47] Z. J. Wang, F. Hohman, and D. H. Chau. Wizmap: Scalable interactive visualization for exploring large machine learning embeddings, 2023. 3

[48] M. Worring, J. Zahálka, S. van den Elzen, M. T. Fischer, and D. A. Keim. A multimedia analytics model for the foundation model era, 2025. 4

[49] Z. Xu, H. Lamba, Q. Ai, J. Tetreault, and A. Jaimes. Counterfactual editing for search result explanation. *CoRR*, abs/2301.10389, 2023. 2

[50] Z. Yang, D. Xue, S. Qian, W. Dong, and C. Xu. Ldre: Llm-based divergent reasoning and ensemble for zero-shot composed image retrieval. In *Proceedings of the 47th International ACM SIGIR Conference on Research and Development in Information Retrieval*, SIGIR '24, p. 80–90. Association for Computing Machinery, New York, NY, USA, 2024. doi: 10.1145/3626772.3657740 2

[51] S. Ye, Z. Sun, G. Wang, L. Guo, Q. Liang, Z. Li, and Y. Liu. Prompt alchemy: Automatic prompt refinement for enhancing code generation, 2025. 9

[52] H. Yin, X. Xu, Z. Wu, J. Zhou, and J. Lu. Sg-nav: Online 3d scene graph prompting for llm-based zero-shot object navigation, 2024. 9

[53] J. Zahálka and M. Worring. Towards interactive, intelligent, and integrated multimedia analytics. In *2014 IEEE Conference on Visual Analytics Science and Technology (VAST)*, pp. 3–12, 2014. doi: 10.1109/VAST.2014.7042476 3

[54] C. Zhao, K. Wang, J. H. Hsiao, and A. B. Chan. Grad eclip: Gradient based visual and textual explanations for clip. *CoRR*, abs/2502.18816, 2025. 3, 6

[55] K. Zhu, Q. Zhao, H. Chen, J. Wang, and X. Xie. Promptbench: A unified library for evaluation of large language models, 2024. 3