

Physical Representation-based Predicate Optimization for a Visual Analytics Database

Michael R. Anderson
University of Michigan
mrande@umich.edu

Michael Cafarella
University of Michigan
michjc@umich.edu

German Ros
Toyota Research Institute
german.ros@tri.global

Thomas F. Wenisch
University of Michigan
twenisch@umich.edu

Abstract—Querying the content of images and video requires expensive content extraction methods. Modern extraction techniques are based on deep convolutional neural networks (CNNs) and can classify objects within images with astounding accuracy. Unfortunately, these methods are slow: processing a single image can take about 10 milliseconds on modern GPU-based hardware. As massive video libraries become ubiquitous, running a content-based query over millions of video frames is prohibitive.

One promising approach to reduce the runtime cost of queries of visual content is to use a hierarchical model, such as a cascade, where simple cases are handled by an inexpensive classifier. Prior work has sought to design cascades that optimize the computational cost of inference by, for example, using smaller CNNs. However, we observe that there are critical factors besides the inference time that dramatically impact the overall query time. Notably, by treating the physical representation of the input image as part of our query optimization—that is, by including image transforms, such as resolution scaling or color-depth reduction, within the cascade—we can optimize data handling costs and enable drastically more efficient classifier cascades.

In this paper, we propose TAHOMA, which generates and evaluates many potential classifier cascades that jointly optimize the CNN architecture and input data representation. Our experiments on a subset of ImageNet show that TAHOMA’s input transformations speed up cascades by up to 35 times. We also find up to a 98x speedup over the ResNet50 classifier with no loss in accuracy, and a 280x speedup if some accuracy is sacrificed.

I. INTRODUCTION

Recent developments in computer vision have made feasible a long-term dream for the database community: a *visual analytics database*, which stores image data and answers user questions about its contents. For example, video frames from a city’s traffic cameras could be used to count cars per minute. Photos uploaded to photo storage web sites could be automatically sorted and tagged based on their contents. The sheer volume and diversity of data captured by cameras opens up myriad analytical query possibilities, if the content hidden behind opaque pixel values can be extracted at scale.

Deep convolutional neural networks (CNNs)—the family of methods used in modern computer vision systems—have enabled huge strides in image understanding in the last few years through tasks like image classification and object detection. For example, the classification error rate in the annual ImageNet Large-Scale Visual Recognition Challenge (ILSVRC) [40] dropped from 25% to 18% in 2012 when a deep CNN was first used [28]. Recent results have lowered the error rate to 2% [20], rivaling or exceeding human performance.

Unfortunately, deep networks pose a considerable computational challenge when deployed in an analytical database system: a model’s inference for a single image can require a lengthy series of large tensor multiplications. For example, YOLOv2, an object detection system designed for speed, requires 8.52 billion operations per single 416x416 pixel image, processing about 67 images per second on a modern GPU [38]. Since GPU hardware is far more expensive than most image sensors, any system requiring a dedicated GPU to process a single video stream will never process more than a fraction of the data produced in real-world, multi-camera applications. Simply, to query huge amounts of image data, we need drastically lower processing costs.

Processing queries over a corpus of image data fits a more general *loop-and-test* pattern that is common to many machine learning tasks: the processor loops over the data, executing an expensive operator on each element to find those satisfying the task’s constraints. In this case, content is extracted from each image by the expensive inference stage of a deep network to determine if the image satisfies a binary predicate specified in a user’s query. While a loop-and-test process can be shortened by processing fewer items overall—using simple sampling or more sophisticated input selection [4]—we focus here on speeding up the *test* phase by reducing the per-image inference cost.

Recent work has reduced inference times for these types of deep learning systems (e.g., [17], [26]). However, we note that all of the visual data system optimizations to date suffer from a critical defect: they concentrate only on computation and ignore the inevitable data-handling costs, such as loading and transformation. Any query optimization method that focuses only on reducing computational load cannot exploit complex *data-centric tradeoffs* that weigh the amount of image data available, the classifier accuracy, and data handling costs.

As an example, consider an optimizer choosing between two image classification models (M_1 and M_2). M_1 accepts a 3-channel, full-color, 224x224 image as input while M_2 accepts a 1-channel, grayscale, 224x224 image. M_1 has fewer convolutional layers and, despite the larger input, requires fewer tensor operations than M_2 , so its inference is faster. M_2 uses less rich data than M_1 , but is nonetheless able to obtain comparable accuracy because of its additional convolutional layers. An optimizer that considers only the inference speed would choose M_1 . However, a data system using M_2 might be faster, as its inputs load in one-third the time of those of M_1 .

Such data-handling tradeoffs are important because visual analytics systems are likely to be architecturally diverse. Some systems may store multiple versions of the same image data (high-res vs. low-res). Others may employ different storage systems (local server vs. cloud vs. in-camera). In some architectures—say, one in which connectivity conditions change, or one in which GPUs are only intermittently available—the highest-payoff query plan may change by the moment. Ignoring data-centric tradeoffs can sacrifice substantial performance. In this work, we propose a framework for handling data-centric tradeoffs when optimizing visual analytical queries.

Technical challenge — We aim to optimize visual analytics queries that might be run in a range of diverse architectures and deployment settings. In this work, we ignore many standard relational query optimization issues, such as query plan rewriting or indexing. We focus exclusively on designing and choosing the CNN-based operator that implements an image-sensitive relational predicate. These operators can be chosen in a manner that trades system throughput against classification accuracy. Making that tradeoff of runtime vs. data quality is an application-specific decision we leave to the user. This paper offers a framework for identifying the best possible operator implementation, subject to a user’s desired tradeoff.

Our approach — One method to speed up an expensive-but-accurate CNN is to replace it with cascades of fast, high precision (but low recall) image classifiers [26], [6], [43]. This is effective but focuses exclusively on computational efficiency. We start similarly, training a large number of specialized candidate binary-classification CNN models by varying not only CNN hyperparameters (as in prior work [26]), but also the representations of the inputs; for example, we build an n -layer CNN for large full-color inputs, another for small full-color inputs, more for grayscale inputs, and so on.

From these core candidate models, we then construct a massive number of classifier cascades. All of these cascades have different initially unknown runtime and accuracy characteristics. Our optimization method efficiently evaluates the cascades’ accuracy using held-out data, and evaluates their runtime characteristics for the system’s current deployment scenario. Finally, it identifies the Pareto-optimal cascades that satisfy the user’s application-specific speed and accuracy constraints.

Organization — After formally describing our problem in Section IV we discuss the following contributions:

- We propose a method for identifying high-quality image predicate implementations, by exploring CNN hyperparameters and varying input data representation (Section V).
- We show the dramatic impact on runtime when a system is running in different deployment scenarios and is aware of the deployment-specific data handling costs (Section VI).
- We prototype our methods in a system called TAHOMA and show that it provides up to a 35x speedup to classifier cascades through input data transformations. TAHOMA also achieves up to a 98x speedup over the ResNet50 image classifier with no accuracy loss (Section VII).

We follow with a discussion of related work in Section VIII.

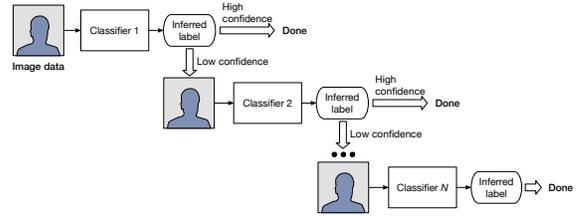


Fig. 1: A multi-level classifier cascade. If the first classifier’s output is uncertain, the input is classified by the second, etc. If reached, the final classifier’s output is accepted as the label.

II. BACKGROUND

In this section, we will briefly give some background on cascade classifiers, since these methods are core building blocks of our work. Numerous systems have used cascades to speed up classifier inference. One of the first was the Viola-Jones face detector [45], which used a series of classifiers based on simple image features to detect faces in subsections of photographs; if any classifier had high confidence in its result, the result was immediately accepted, avoiding the use of further classifiers.

Figure 1 illustrates the general cascade process: an image is input into a classifier, whose output is accepted if it has high confidence. Otherwise, a second classifier is used, where again, a low-confidence result will send the input to a third, and so on. If the final level of the cascade is reached, its output is accepted. Ideally, the initial levels of a cascade are fast with high precision, though they may suffer from low recall. In face detection, for example, if most images have no human faces at all, the first classifier may quickly eliminate most cases. Only the few cases containing a face will be processed by the remainder of the cascade, at correspondingly higher cost.

Recently, cascades have been used to accelerate the relatively slow inference speeds of deep neural networks [26], [6], [43]. Our system takes these techniques further by exploiting the representations of the inputs and finding optimal cascades for a user’s deployment scenario.

III. DESIGN CONSIDERATIONS

Several key questions are important to TAHOMA’s design. We touch on these issues below.

Issue 1: Object detection vs. image classification — In general terms, an object detector (e.g., YOLOv2 [38]) finds the location of particular object classes within an image, while an image classifier (e.g., ResNet50 [16]) identifies the overall contents of an image as one of a particular set of classes. When implementing binary `contains-object` predicates in a visual database application, the extra architectural and computational complexity of object detectors is unnecessary. Our cascades therefore use small and fast CNN image classifiers. Applications requiring object location within images, however, could build cascades from small and fast object detectors.

Issue 2: Online vs. offline classification — If a user’s corpus is small or slow-growing enough to allow for offline classification of the entire dataset with available resources, this paper’s techniques are unnecessary. Also, our techniques may be unneeded if query predicates are fixed and new ones

are unlikely to be introduced. In such cases, the user could materialize the classification results for each image on ingest and store them in a standard database for future queries.

However, there are many exciting real-world applications where assumptions of a small or slowly growing dataset or stable query predicates are unrealistic. Consider, for example, a website for photo storage: Flickr has reported that days of heavy usage can see 25 million photo uploads [1]. Facebook had 350 million daily photo uploads as of 2013 [10]. Video applications can be even more extreme. A single self-driving car can have over a dozen cameras, each gathering 30 frames per second for many hours on a daily basis; a fleet of such cars can generate a huge amount of data.

Further, many applications require retrospective exploration and analysis, where the query predicates may not be known in advance. For example, a self-driving car engineer may wish to find historical examples of a new failure case, requiring the training of a new model to be run over an existing corpus. Or consider a police investigation reviewing thousands of hours of surveillance camera footage to find a local delivery van with a unique logo: a trained object detector would be ideal, but such a specific model is unlikely to already exist. In general, it is unlikely that all possible query predicates can be enumerated in advance for a visual analytics application.

Issue 3: Training costs — A deep CNN image classifier, such as ResNet50 [16], can take days to fully train, due to the model’s architectural complexity and the huge training set needed for such a complex model. Such a burden is incompatible with our use case, since requiring this amount of time to install a new predicate in our system is unreasonable. Thankfully, our simple binary predicates typically do not require huge classifiers. The majority of our cascades consist solely of small, specialized classifiers, which train in just minutes.

Further, in cases where a deep network may be needed, training it from scratch is likely to be unnecessary. It is common practice to fine-tune pre-trained deep networks to a particular task [24]. Generally, when fine tuning, most of the deep network is frozen and only the last several layers are modified and retrained to a new task, taking advantage of already-learned features from a similar problem domain. In our experiments, we fine tuned ResNet50 for binary classification tasks using a modern GPU in only 2 to 4 hours.

Issue 4: Deployment scenarios — A visual analytics system may be deployed in a variety of scenarios, requiring accounting for differing data handling costs, in addition to classifier inference costs. Consider the following example scenarios, also later used in our experiments:

- **ARCHIVE** – In this situation, a large archival corpus of historical image data is stored on local drives. Each image must first be loaded from the drive and then transformed into an appropriate input format for the classifier.
- **ONGOING** – Here, video is continually ingested from its source into a datacenter-based query system, where it is transformed into appropriate representations that are stored on SSD for later queries. Because this data is

TABLE I: Frequently used notation.

Notation	Definition
$\mathcal{J} = (I_1, \dots, I_n)$	Image data corpus
$T_i = (t_{i1}, \dots, t_{im})$	Content tuple for image I_i
K	Image classifier, s.t. $K(I_i) = t_{ij}$
$\mathcal{M} = (M_1, \dots, M_m)$	Basic classification models
$\mathcal{A} = (A_1, \dots, A_{n_a})$	Model architecture specifications
$\mathcal{F} = (F_1, \dots, F_{n_f})$	Input transformation functions
$C = (M_1, \dots, M_n)$	Cascade of n models belonging to \mathcal{M}
$\mathcal{C} = (C_1, \dots, C_k)$	Collection of cascades for a binary query
p_{low}, p_{high}	Cascade model decision thresholds

transformed as it is acquired, only the cost of loading the representations from disk are considered at query time.

- **CAMERA** – If compute nodes are at the edge of the network (e.g., connected to surveillance cameras), the images can be directly provided to the classifiers. Only the image transformation costs must be considered, since transfer costs from camera to memory are negligible.

We show in Section VII-A that scenario-awareness while evaluating of the accuracy and speed of classification systems can lead to a large practical increase in a system’s throughput.

IV. DEFINITIONS AND NOTATION

This section formalizes the particular image classification problem addressed here. Table I lists frequently used notation.

Content-based Queries — A query system that operates over images can perform queries over two main types of information: image metadata (e.g., GPS or acceleration information that may accompany frames from dashcams) and content extracted from images themselves (e.g., the image contains a bicycle). Metadata queries are easily handled by existing methods, so we focus this work on processing *content-based queries*.

Definition 1 (Content-based query): Given a corpus of image data \mathcal{J} containing images I_1, \dots, I_n , the tuples $T_i = (t_{i1}, \dots, t_{im})$ represent the visual contents for each image I_i , where each element t_{ij} in this tuple represents a content object present in image I_i . A content-based query is constructed of predicates that can be evaluated with the elements of T_i .

We restrict content-based queries to *binary queries*, which can be combined into a full query over all T_j tuples.

Definition 2 (Binary query): Given an image’s content tuple T_i , a binary query involves a single `contains-object` predicate evaluated with a single tuple element, t_{ij} . A binary query thus asks if image I_i contains object t_{ij} .

Despite this restriction, our envisioned system will support complicated queries that can be rewritten as a combinations of metadata and binary query predicates. For example, the query “Find images from Detroit containing a bicycle” can be decomposed into a metadata predicate (`location = ‘Detroit’`) and a binary query predicate (`contains_object(bicycle)`). Our focus is on choosing the best classifiers to implement a given `contains_object` predicate. While further query optimization could be done considering multiple binary predicates in concert, we leave that for future work and here concentrate on optimizing single predicates.

Image Classification — The content tuples are generally not available upon image acquisition and must be extracted from the image. In this work, we focus on extraction via *image classification*. In particular, we focus on queries where classification has not been performed in advance, so that classification must be performed as part of query execution.

Definition 3 (Image classification): Given a corpus of image data \mathcal{J} , for each image $I_i \in \mathcal{J}$, we wish to generate a binary label L using a classifier K , such that $L_i = K(I_i)$. The label L_i corresponds to a member $t_{i,j}$ of content tuple T_i .

The output of a classifier model can be thought of as a virtual column in a relation describing the content objects in images. For example, processing the query predicate `contains_object(bicycle)` would populate the bicycle column of this relation with the output of K_{bicycle} . The classifier K could be generated by a range of methods, including *basic models* like logistic regression or deep CNNs, or may be a collection of basic models, such as *classifier cascades*.

Definition 4 (Basic model): A basic model M implements a classification method that accepts image I as input and outputs a binary classification result. Our system generates a large set of such CNN-based models $\mathcal{M} = (M_1, \dots, M_m)$.

Two factors parameterize a model M : *model architecture specification* A and *input transformation function* F .

Definition 5 (Model architecture specification): The internal architecture of a model M is specified by A . For the CNNs used in our system, A_m describes network hyperparameters, such as the number and size of layers. $\mathcal{A} = (A_1, \dots, A_{n_a})$ gives all potential architectural options for models in \mathcal{M} .

Definition 6 (Input transformations): Before the classification of an image I by model M , the raw image data is processed by an image transformation function F , such that input image I is transformed into output image I' . Such a function may perform one or more operations such as resizing, normalizing, or reducing color depth. The set $\mathcal{F} = (F_1, \dots, F_{n_f})$ gives all functions available to pre-process image data for models in \mathcal{M} .

We use the cross product of \mathcal{F} and \mathcal{A} as the model design space, resulting, in practice, in several hundred individual models in \mathcal{M} for each binary query. A goal of this work is to determine which models are most suitable for a user’s accuracy and runtime constraints and current deployment scenario.

Classifier cascades — Classification models can be aggregated into collections or ensembles to improve either accuracy or speed. One such method designed to improve classification speed is the *classifier cascade* [45].

Definition 7 (Classifier cascade): A classifier cascade $C = (M_1, \dots, M_n)$ is a list of n basic models with probabilistically interpretable output. The models are run in series: image I_i is classified by M_1 , and if the output is between two given decision thresholds, p_{low} and p_{high} , it is *uncertain*. If so, I_i is then classified by M_2 . Otherwise, the cascade is stopped and M_1 ’s output is accepted as the label of I_i . This continues to the final classifier M_n , whose output is always accepted.

Given a set of classification models \mathcal{M} , we can construct a large set of cascades $\mathcal{C} = (C_1, \dots, C_k)$ with up to n levels each, to be evaluated in terms of accuracy and throughput.

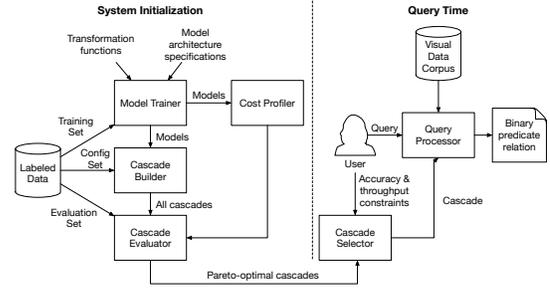


Fig. 2: TAHOMA architecture

Model evaluation — The quality of classifier—either a model M or a cascade K —is given by its *accuracy* and its *throughput*. Accuracy gives the fraction of labels produced by M that are correct. Throughput is the number of classifications per unit time and measures how fast the model’s relation is populated.

For a set of classifiers, we can find a subset that is Pareto-optimal over these two criteria. That is, there is a subset that is non-dominated in terms of accuracy and throughput, from which users can select a classifier to meet application needs.

Problem statement — With the preceding definitions, we can formally describe the problem addressed in this paper thusly:

For an image corpus \mathcal{J} and a set of binary classification models $\mathcal{M} = (M_1, \dots, M_m)$ parameterized by architectural specifications $\mathcal{A} = (A_1, \dots, A_{n_a})$ and input transformation functions $\mathcal{F} = (F_1, \dots, F_{n_f})$, find the set of classifier cascades $\mathcal{C} = (C_1, \dots, C_k)$ constructed from models in \mathcal{M} that are Pareto-optimal in terms of accuracy and throughput over \mathcal{J} .

V. CASCADE METHODOLOGY

In this section, we give an overview of system architecture and discuss the details of our models and cascades.

A. System Architecture

Figure 2 sketches out TAHOMA’s architecture. TAHOMA has two main modes of operation: system initialization and query execution. During system initialization, model repository is prepared for each binary predicate, which requires a set of labeled data. This dataset is small compared to what is generally used to train deep CNNs: per binary predicate, TAHOMA requires 3,000–4,000 labeled images, with equal numbers of positive and negative examples. The labeled data is split into three sets for training, configuration, and evaluation. Training set $\mathcal{J}_{\text{train}}$, transformation functions \mathcal{F} , and architecture specifications \mathcal{A} are provided as input to the model trainer.

For a given binary predicate, a set of models \mathcal{M} (each implementing the `contains_object` operator) is trained and provided to the cost profiler and to the cascade builder. The cost profiler measures the throughput of each model in the current deployment scenario (see Section VI). The cascade builder constructs TAHOMA’s cascade set \mathcal{C} , using all possible combinations of size n of the models in \mathcal{M} and the configuration set $\mathcal{J}_{\text{config}}$ (see Section V-C). Using the evaluation set $\mathcal{J}_{\text{eval}}$, the cascade evaluator measures each cascade’s accuracy and throughput (see Section V-E). With this, the system determines the set of Pareto-optimal cascades for use at query time.

Similar to how approximate query systems like BlinkDB [3] and VerdictDB [34] allow users to specify approximation constraints as part of their queries, a TAHOMA user provides their constraints on accuracy (U_{acc}) and throughput (U_{thru}) at query time (in the form of the highest tolerable loss in either of those parameters). The cascade selector chooses which of the Pareto-optimal cascades best suits the user’s desired tradeoff. For example, the user may wish to maximize throughput as long as the resulting cascade does not suffer more than a 5% loss in accuracy over the most accurate cascade available. The user would set $U_{\text{acc}} = 0.05$ and provide no constraint for U_{thru} . The system would select the cascade from the set of Pareto-optimal cascades has an accuracy closest to (but not below) 95% of that of the most accurate cascade. Because this is a Pareto-optimal choice, there will be no faster cascades at that (or a higher) accuracy level. The selected cascade processes the data in the corpus, extracting the notional relation for the binary predicate in the user’s query.

Integration considerations — Because our goal was to explore the optimization methods discussed in this paper, we implemented TAHOMA as a standalone query system. However, we believe future deployments of TAHOMA’s ideas will likely be embodied in RDBMS software. The execution of a `contains_object` operator is analogous to that of a user-defined function (UDF) in a database such as PostgreSQL, and could be wrapped in the RDBMS `CREATE FUNCTION` statement. RDBMS query optimizers could leverage additional metadata relations, such as image location and capture date, to reduce the number of expensive TAHOMA UDFs calls for a specific query. Further, UDF output could be stored as a partially materialized table, enabling further query optimization.

TAHOMA’s initialization process is run at the installation of each new binary predicate. During this, the profiled speeds could be used to inform the RDBMS query optimizer of the execution cost of the UDF (like the PostgreSQL’s `COST` parameter for `CREATE FUNCTION`, if each Pareto-optimal cascade was implemented as a separate UDF). While indexing or materializing the output of the UDFs at system installation is not practical in our envisioned deployment scenarios, database triggers could be used to execute the TAHOMA UDFs over newly ingested data after system initialization to pre-materialize the output for future queries. In such situations, slower processing may be tolerated for more accurate results, allowing a different Pareto-optimal cascade choice than at query time.

B. Building Models

With TAHOMA, we create a huge number of cascades \mathcal{C} by first training a large number of individual classification models \mathcal{M} . We build the collection of models in two ways: by varying the internal architecture of our CNN-based classifiers with our model architecture specifications \mathcal{A} and by transforming the input images using the input transformation functions \mathcal{F} .

Model architecture variations — TAHOMA uses convolutional neural networks for its models. Goodfellow et al. provide in-depth discussion on CNNs and deep learning [15]. Our

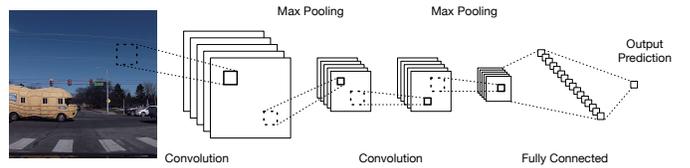


Fig. 3: CNN architecture used by TAHOMA. The number of layers, as well as the number of nodes in each layer is varied as part of the model architecture specifications \mathcal{A} .

CNNs follow the basic architectural pattern shown in Figure 3. Input values are fed into one or more layers of convolutional nodes. Each convolutional layer is followed by a max pooling layer, connected by rectified linear activations (ReLU). The final convolutional layer feeds into a fully connected ReLU layer. A sigmoid output node provides the inferred label. A key point is that our CNNs are small (and thus fast), typically having only one to four convolutional layers. When creating models, we vary these architectural details according to \mathcal{A} ; for our experimental settings of \mathcal{A} , see Section VII-A.

Input transformations — We also vary the physical representation of the input to each model. The set of input transformation functions \mathcal{F} comprises functions that perform one or more image processing operations, such as resolution scaling and color channel modification. These types of transformations are useful for building fast, small models: reducing image size and color depth can greatly reduce the number of model input values, directly reducing the size of the CNN’s tensor operations. For our experiments, we scaled the image resolution (30x30, 60x60, 120x120, and 224x224 pixels), and for each image size, we used five different color variations (full 3-channel color, each of the individual red, green, and blue color channels, and single-channel grayscale).

The design space defined by these model architecture variations and input transformation functions result in hundreds of different model configurations (360 in our experiments). Once each model is trained on a labeled subset of \mathcal{J} , we can compose the models into cascades. Training can take less than a minute for the smallest networks (one convolutional layer with few nodes) with the smallest inputs (30x30 pixels, 1 color channel) to nearly an hour for the largest. Overall, training 360 models for a single binary predicate requires about 12 hours when done serially on an NVIDIA Tesla K80 GPU. Training is parallelizable, so this cost can be greatly reduced in practice.

C. Computing Decision Thresholds

Each model in \mathcal{M} provides a probabilistic output for a binary classification problem, a real number ranging from 0 to 1. Each model has a pair of decision thresholds, p_{low} and p_{high} , which determine whether the model’s labelling decision should be trusted. If the output $o \leq p_{\text{low}}$ or $o \geq p_{\text{high}}$, the model’s output is accepted as the output of the cascade. If $p_{\text{low}} < o < p_{\text{high}}$, then we consider the model’s output to be uncertain and reclassify the image with the next model in the cascade.

These thresholds are chosen on a per-model basis, such that the precision of classification results with $o \leq p_{\text{low}}$ or $o \geq p_{\text{high}}$

matches a predefined constraint while recall is maximized. Using a small configuration dataset distinct from the training set, the thresholds are determined via a grid search that sweeps through the potential thresholds to find those that provide a precision value greater than or equal to the target precision and selects the thresholds from those that maximize recall.

D. Constructing Cascades

Each model has its decision thresholds determined independently, rather than in the context of a specific cascade. This assumption of independence allows us to quickly instantiate and evaluate the millions of possible multi-level cascades that can be constructed from the models in \mathcal{M} .

To determine the accuracy and throughput of our cascades, we first classify a set of labeled images J_{eval} with each model in \mathcal{M} . (The images in J_{eval} are distinct from those for training and determining decision thresholds so that the resulting accuracy measurements are not the product of overfitting.) Since the cascades in \mathcal{C} comprise combinations of the models in \mathcal{M} , the above evaluation need only be done once per model (360 times in our experiments) and not once per cascade (1.3 million).

Ensuring the independence of both model evaluations and decision thresholds (as described in Section V-C) enables extremely fast evaluation of cascades: our evaluation required just over one minute to determine the accuracy and throughput values for 1.3 million cascades. As such, once models have been trained and classified, the selection of a cascade can be part of query planning at query execution time and can thus incorporate query-specific performance criteria (e.g., which storage devices are providing input images).

Using the pre-computed classification results for J_{eval} for each model in the cascade, the cascade execution is simulated to obtain its label predictions for J_{eval} . The cascade’s accuracy is then computed by comparing with true labels for J_{eval} .

E. Evaluating Cascade Sets

Each cascade in \mathcal{C} can be displayed on a plot of accuracy versus throughput, as shown in Figure 4. Those with the best tradeoffs between accuracy and throughput belong to the Pareto frontier. Points on the Pareto frontier are those not dominated by any other points [33]. Computing the Pareto frontier over two attributes, as we do, is $O(n \log n)$ in the number of points [30]. Here, n is the number of cascades, and while this numbers in the millions, this is a fast computation, as just described.

VI. DATA HANDLING COSTS

An often overlooked part of image classification is the cost of loading and preparing the data prior to inference. Of the rare projects in the computer vision field that report inference speeds (e.g., the YOLO family of object detectors [37], [38]), none report video decoding or image loading time. Likewise, the NoScope video query system [26] explicitly ignores these costs, claiming that GPU-based decoding is so fast as to be negligible or that decoding might be avoided altogether by obtaining raw video frames directly from the generating camera sensor.

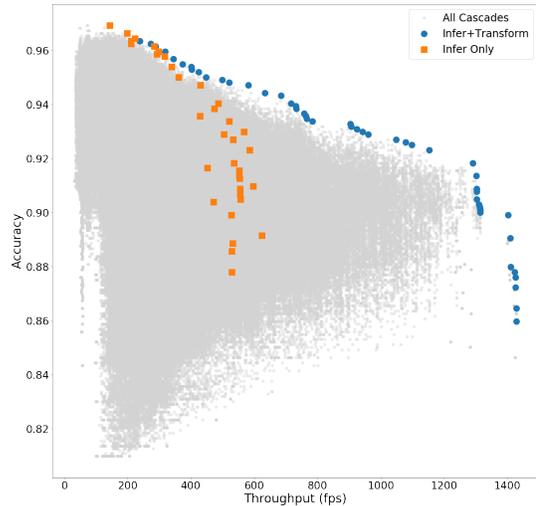


Fig. 4: Cascades (gray) and Pareto frontier (blue) for an example deployment scenario, compared to the Pareto frontier for a scenario only considering inference costs (orange).

This last point hints at why image loading or decoding costs *should* be included when evaluating these systems: deployment scenarios for a given query system may differ drastically. Some may use top-of-the-line GPUs for video decoding, others may store video frames as individual image files on disk, while some may require video to be transported over a network prior to processing. Further, if multiple classification models with a variety of input representations might be used for classification, data handling costs—including preprocessing—must be included in throughput evaluations when deciding which classifier will be used at query time.

More concretely, for accurate model comparisons to be made, the throughput must be measured as the reciprocal of the average classification time, t_{classify} , defined as follows:

$$t_{\text{classify}} = t_{\text{load}} + t_{\text{transform}} + t_{\text{infer}}$$

where each time t above is the average recorded over some typical set of input images, measured on the deployed system.

Our experiments in Section VII-D demonstrate that choosing among cascades with incorrect cost assumptions can lead to a large decrease in throughput. This throughput loss can be seen in Figure 4. The gray points depict all possible cascades for an example binary predicate (e.g., `contains-object (semitruck)`) in a deployment scenario where image loading costs are negligible (full size raw data is already present in memory), but image preprocessing costs are incurred to transform the images into the appropriate resolutions and color representations for each model. The points on the Pareto frontier (shown in blue) represent the cascades that present the best tradeoffs between accuracy and throughput under this deployment scenario. The orange points show the cascades that would be on the Pareto frontier for this binary predicate if only inference costs were considered (pre-processed images are already present in memory). *If the preprocessing costs were not considered, the query throughput would be far below its potential for most accuracy levels.*

VII. EXPERIMENTS

We have implemented TAHOMA as a prototype system, and in this section, we discuss our experiments that compare TAHOMA’s performance against existing baseline methods, investigate the effects of data handling costs of different deployment scenarios, evaluate the effects of input transformations, and analyze the effects of increasing cascade depth.

A. Experiment Setup

To evaluate the methods and components used in TAHOMA, we designed a series of experiments with the following setup.

Binary predicates — We evaluated TAHOMA’s performance over a set of 10 queries with a single `contains-object` binary predicate, randomly chosen from the 1,000 categories in the ImageNet dataset [40]. Each query is of the form “SELECT * FROM images WHERE `contains-object(category)`”, where `category` is one of those shown in Table II. The ImageNet dataset provides about 1,200 training images for each category. We held out 200 of each category’s images as a validation set. The remaining images were labeled as positive examples. We then selected a matching number of random images drawn from the remaining 999 ImageNet categories and labeled them as negative examples. Because these training sets were relatively small, we followed common data augmentation practice by creating a copy of each image that was flipped left-to-right, doubling the amount of training data.

We performed similar labeling for the validation set. For performance evaluation, we collected an additional dataset for each category using the image search functionality of the Google and Bing search engines, resulting in roughly 500 positive images per category. To find negative examples, we first created a large collection of images by performing image searches on Google and Bing for each of the remaining 999 ImageNet categories. From that, we randomly chose our negative examples to match the number of positive ones.

Cascade configurations — We used the Keras [8] deep learning library running TensorFlow[2] to train and execute the CNNs used in our cascades. We varied several network architecture hyperparameters to create a range of models for each binary predicate: the number of convolutional layers (1, 2, 4), the number of convolutional nodes in each layer (16, 32), and the number of nodes in the final dense layer (16, 32, 64). These hyperparameters provide a range of reasonable options similar to those used in other systems (e.g., [26]). We also varied the size of the input images (30x30, 60x60, 120x120, 224x224). For each input size, we used five different image representations: full 3-channel color, each of the individual red, green, and blue color channels, and single-channel grayscale. In all, we constructed 360 simple CNNs for each binary predicate.

We also included in our pool of classifiers a fine-tuned implementation of ResNet50 [16] that had been pre-trained on ImageNet. Fine-tuning was done using standard techniques; the final 1000-class classification layer was replaced with 64-node ReLU dense layer, followed by a 2-node softmax layer for the

TABLE II: ImageNet categories randomly selected for use as our experimental binary predicates.

Predicate	ImageNet ID	Query	ImageNet ID
1. acorn	n12267677	6. ferret	n02443484
2. amphibian	n02704792	7. komondor	n02105505
3. cloak	n03045698	8. pinwheel	n03944341
4. coho	n02536864	9. scorpion	n01770393
5. fence	n03930313	10. wallet	n04548362

output of the binary prediction. These layers were trained using the same training set as the smaller, specialized classifiers.

For each classifier, we calibrated its decision thresholds using the validation dataset for five precision settings: 0.91, 0.93, 0.95, 0.97, and 0.99. We use each classifier variation to construct cascades of one and two levels, as well as three-level cascades with ResNet50 as the final layer (see Section VII-F), resulting in 1,301,405 possible cascades per predicate.

Deployment Scenarios — The data handling and preprocessing costs of particular deployment scenarios can have a large effect on the optimal choices in classifier cascades. To demonstrate this, we analyzed our cascades under four different scenarios, corresponding to those in Section III:

- **INFER ONLY** — This scenario ignores data handling and transformation costs—only inference costs (i.e., only the time required to evaluate the CNN) are considered when computing throughput, a practice commonly used in computer vision literature. However, as we show, the fastest inference often does not imply the fastest end-to-end query performance in practical deployments.
- **ARCHIVE** — This scenario includes the cost of loading a full-size image off an SSD hard disk, as well as the cost of resizing that image to the appropriate input size for a given classifier, as might be done when querying an existing corpus of archived image or video data.
- **ONGOING** — This scenario corresponds to deployments where images are resized on ingest before saving to disk. Load times are smaller, since the full-sized image is not loaded from disk if not needed for a particular classification. This scenario may occur when setting up a data collection system in tandem with a query system; proper image sizes for object detectors are known and initial transform costs can amortized over many queries.
- **CAMERA** — This scenario only includes the computation costs of resizing the images, as in deployments where loading costs are negligible (e.g., images are loaded to memory directly from a connected camera sensor).

Data handling costs for the above scenarios only occur once for a given input: if a cascade includes two classifiers that use, for example, a 30x30 pixel red channel input, the costs to create that input are incurred only once per image.

Cascade evaluation — To compare two sets of cascades in terms of throughput, we compute the area to the left of the curve (ALC) created by those points in this plot over a given accuracy range. Because a Pareto frontier is a collection of points and not a curve, we interpolate the curve as a step function. Dividing

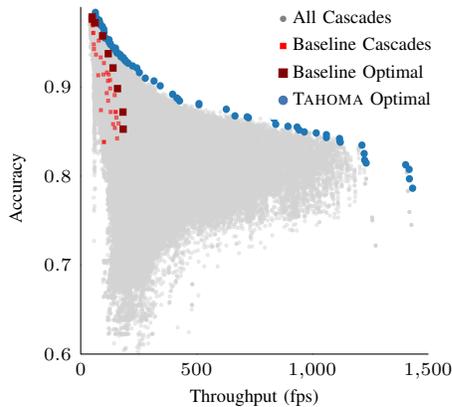


Fig. 5: A comparison of the cascade space of TAHOMA (gray) compared to that of our Baseline cascades (red). TAHOMA’s Pareto-optimal points are in blue. This example uses our komondor binary predicate under the CAMERA cost model.

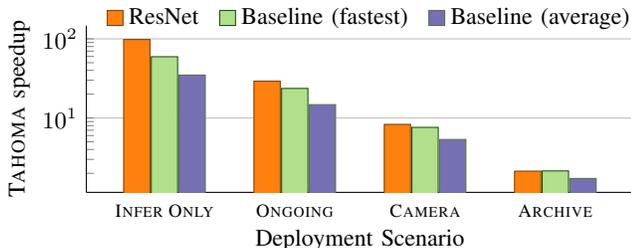


Fig. 6: Average speedup values of TAHOMA over baselines. *ResNet50* and *Baseline (fastest)* comparisons use the optimal cascade with the nearest higher accuracy to *ResNet50* and the fastest *Baseline* cascade, respectively. *Baseline (average)* shows average speedups over the *Baseline* accuracy range.

the ALC by the size of the accuracy range give the *average throughput* for cascades in the Pareto frontier. Dividing the ALC of one frontier by another gives the *speedup* of the former over the latter. For fair comparisons, we use the accuracy range for the full set of cascades for each configuration and choose the smallest said range. In some cases, we compute the ALC for a Pareto frontier’s cascades in a different cost context. These cascades are no longer a strict Pareto frontier, but we can still compute ALC for comparisons. When comparing TAHOMA to a single classifier, such as against our *ResNet50* baseline system, we choose the optimal cascade whose accuracy is both higher and closest to the accuracy of the single classifier.

Hardware — We used an Amazon EC2 p2.xlarge instance with an NVIDIA Tesla K80 GPU to do training, inference, and throughput measurements for all CNNs. For operations suited to parallel CPU processing—such as finding the Pareto frontiers for our cascade sets—we used a 32-core (2.8GHz) Opteron 6100 server with 512GB RAM.

B. Comparison Against Baselines

Overview — We compare TAHOMA against a fine-tuned, pre-trained *ResNet50* implementation [16], as well as a set of non-optimized cascades, which comprise a subset of TAHOMA’s design space. These two-level cascades terminate in a full-cost

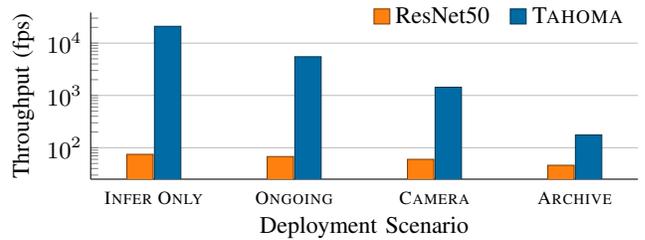


Fig. 7: Throughput of TAHOMA and *ResNet50* of fastest cascades for each cost model, averaged over 10 binary predicates.

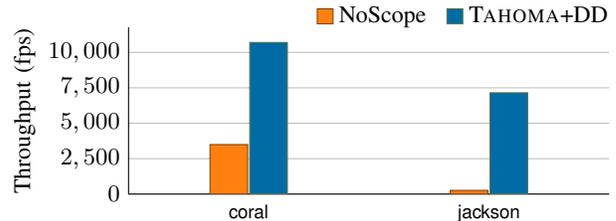


Fig. 8: Comparison with NoScope. TAHOMA+DD is TAHOMA with a simulated NoScope-style difference detector.

classifier (i.e., *ResNet50*) and use full-color 224x224 images as input. These are similar to design to CNN-based cascades in previous work [26]. An illustration of the difference in the design spaces is shown in Figure 5, with TAHOMA’s available cascade options being markedly larger due both the use of data transformations on the inputs and the additional cascade depth.

Throughput gains — Figure 6 shows the performance of TAHOMA compared to our baselines for the four deployment scenarios from Section VII-A. When only considering inference speed (i.e., INFER ONLY), TAHOMA yielded a 98x speedup over using a fine-tuned *ResNet50* classifier alone, averaged over our 10 predicates. TAHOMA showed a 35x average speedup over the accuracy range provided by the *Baseline* cascades. At the accuracy level provided by the fastest *Baseline* cascade for each predicate, TAHOMA yielded a 59x speedup on average. For the other scenarios, data handling overheads reduce the speedup gains. Nevertheless, TAHOMA achieves substantial speedup in all scenarios, and even the most expensive scenario that requires costly loading and transformation costs (ARCHIVE) shows a nearly 2x speedup versus both *ResNet50* and *Baseline*.

If speed is the priority, TAHOMA allows a user to trade accuracy for a large throughput boost. Figure 7 shows TAHOMA’s fastest optimal cascade compared to *ResNet50*. Across all predicates, the fastest cascades were not true cascades at all: they comprised a single specialized classifier with adequate accuracy and high throughput. In the INFER ONLY scenario, TAHOMA achieved an average throughput of 20,926 frames per second—280 times faster than our fine-tuned *ResNet50* models, which had an average throughput of about 75 frames per second. The more realistic ONGOING scenario still achieves 5484 frames per second—an 81x speedup. These large speedups come at a price, though: under INFER ONLY, the TAHOMA’s fastest cascade was on average 12% less accurate than *ResNet50*. Of course, as Figure 5 shows, the optimal cascades provide a rich space of throughput and accuracy tradeoffs, so users can find the right balance for their needs.

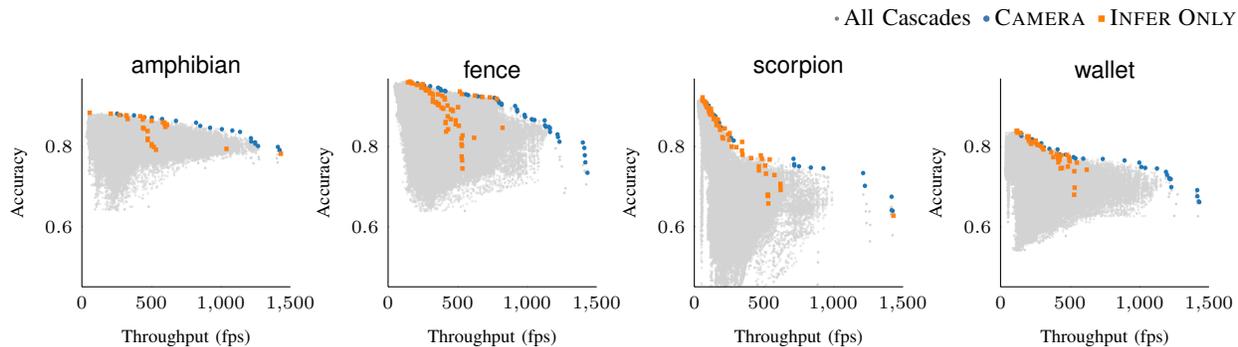


Fig. 9: Pareto frontiers for several of our binary predicates, under the CAMERA deployment scenario (blue), compared to cascades that would be in the Pareto frontier for the INFER ONLY scenario (orange). Optimal cascades under one deployment scenario can significantly underperform those in another, so consideration of realistic deployment scenario costs is crucial.

C. Comparison with NoScope

Because NoScope [26] is the existing system most closely aligned with our work, we ran experiments to directly compare the two systems. For these experiments, we used the code and datasets (coral and jackson) provided by the NoScope authors¹. The other datasets presented in the NoScope paper were not publicly available. We used the default parameters provided in the NoScope code for each dataset and report results for both systems with a target precision of 0.95 used to select cascade thresholds. YOLOv2 [38] was used as the final, expensive classifier for both systems. Both NoScope and TAHOMA were run on AWS p2.xlarge instances. Note that CPU and GPU specifications differ between our NoScope installation and the one in the NoScope paper, so raw performance numbers differ between our experiments and theirs.

To compare NoScope and TAHOMA on an equal footing, we implemented TAHOMA+DD, which is TAHOMA with a simulated difference detector equivalent to that used by NoScope. The difference detector measures the similarity between the current frame and previously seen ones and reuses previous results if the compared frames meet a similarity threshold. This mechanism is orthogonal to our work and increase NoScope’s throughput by avoiding many classifier executions. To create TAHOMA+DD, we recorded frame similarity using NoScope’s difference detector and reused TAHOMA’s results for frames meeting NoScope’s threshold instead of classifying them.

Additionally, both systems used basic frame skipping, only processing one of every 30 frames. The results shown here include only those frames actively processed by each system, not those skipped this way. TAHOMA+DD results are measured in the INFER ONLY deployment scenario, which matches NoScope’s throughput measurements. TAHOMA+DD results use the Pareto-optimal cascade with the closest but higher accuracy level to that of the NoScope for each dataset.

Figure 8 compares the throughput for NoScope and TAHOMA+DD for the two public NoScope datasets. For both datasets, TAHOMA+DD significantly outperformed NoScope. On coral, TAHOMA+DD system reached a throughput of

10,700 fps, while NoScope’s throughput was 3494 fps, giving TAHOMA+DD a 3.1x speedup over NoScope. On jackson, TAHOMA+DD system reached a throughput of 7,150 fps, while NoScope’s throughput was 260 fps, giving TAHOMA+DD a 27.5x speedup over NoScope².

D. Deployment Scenario Awareness

Figure 9 shows (in blue) the Pareto frontier of all classifier cascades for the CAMERA cost model for several of our binary predicates. Additionally, the cascades that would be Pareto-optimal for each query under the INFER ONLY model are shown in orange. These orange points form a non-convex curve, since they are *not a Pareto frontier* under the depicted CAMERA cost model. Each cascade may be impacted differently by loading and transformation costs, so their throughputs can change relative to one another in different deployment scenarios. With few exceptions, the optimal cascades under CAMERA are different than the INFER ONLY ones. It is clear that if the data handling costs of a scenario like CAMERA were ignored and the “optimal” cascades were chosen only considering inference costs, considerable throughput gains would be missed.

Table III shows the difference in throughput in our deployment scenarios when cascades are chosen in a scenario-oblivious way (i.e., when only inference costs are considered, as in INFER ONLY) versus when the cascades are chosen taking scenario costs into consideration. Because TAHOMA provides a tradeoff between accuracy and throughput, we show results for four different levels of permissible accuracy loss. A user, for example, may decide that a 5% decrease in accuracy is acceptable in order to process images faster. Then, in the CAMERA scenario, the system’s throughput would increase by 59.5% if cascades were chosen taking data handling costs into consideration, instead of being oblivious to these costs and only considering the classifier’s inference.

² As may be apparent from the results, the coral dataset was a much simpler classification task than jackson, with far more reused results from the difference detector for coral (25.2% reused) than for jackson (3.8% reused). NoScope used the expensive YOLOv2 model for a significant number of frames on jackson, as well, leading to its slow performance. TAHOMA+DD’s much larger cascade design space allowed it to find an accurate cascade that was able to avoid calling YOLOv2 for all but a few frames.

¹<https://github.com/stanford-futuredata/noscope>

TABLE III: Throughputs for various deployment scenarios when the cascade choices chosen in either oblivious or aware of scenario data handling costs. Here, permissible accuracy loss indicates how much accuracy the user is willing to trade for an increase in throughput. Scenario awareness can lead to significant throughput increases, shown in parentheses.

Permissible accuracy loss	Scenario: ARCHIVE		Scenario: CAMERA		Scenario: ONGOING	
	Oblivious	Aware	Oblivious	Aware	Oblivious	Aware
0% loss	57.5 fps	58.3 fps (+1.4%)	107.1 fps	107.1 fps (+0.0%)	111.9 fps	111.9 fps (+0.0%)
2% loss	85.1 fps	91.1 fps (+7.1%)	267.5 fps	324.6 fps (+21.3%)	985.2 fps	985.3 fps (+0.0%)
5% loss	103.1 fps	117.1 fps (+13.5%)	344.7 fps	549.9 fps (+59.5%)	1938.7 fps	2000.8 fps (+3.2%)
10% loss	130.6 fps	142.0 fps (+8.7%)	568.0 fps	806.8 fps (+42.0%)	3669.1 fps	3669.1 fps (+0.0%)

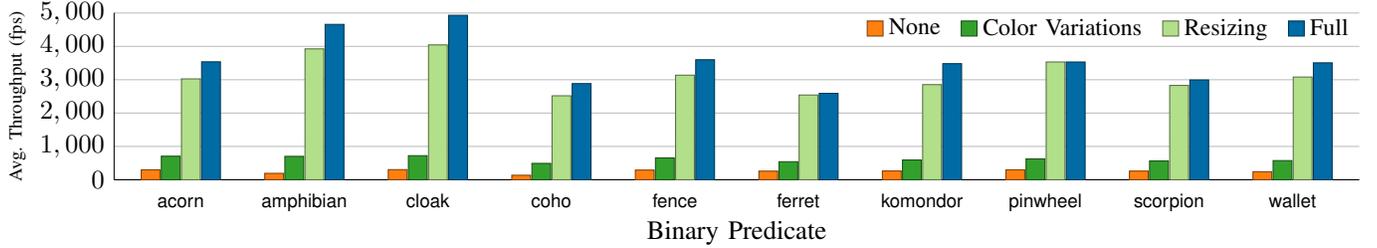


Fig. 10: Average throughput of optimal cascades for cascade sets that use different input transformations.

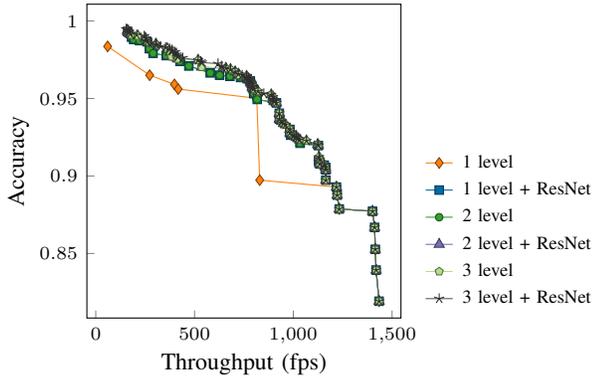


Fig. 11: Evolution of Pareto frontier as cascades depth increases, shown for our fence predicate in the CAMERA scenario. Other predicates and scenarios showed similar results. As cascades get deeper, Pareto frontier improvements become negligible.

E. Analysis of Input Transformations

TAHOMA uses several different input transformations to expand the space of simple classifiers used to construct cascades. To see how these affect TAHOMA’s performance, we constructed four cascade sets that used varying subsets of the transformations: *None*, which used no input transformations (i.e., all inputs are 224x224 three-color-channel images); *Color Variations*, which used only the transformations that extract the color channels or create grayscale images; *Resizing*, which used only transformations that reduce resolution; and *Full*, which used the full set of transforms included in TAHOMA.

Figure 10 shows the average throughput for each cascade set for each binary predicate, computed using the ALC method described in Section VII-A. We computed these values over the accuracy range of the *Full* cascade set for each predicate. Image resizing operations by far have the largest impact on throughput, giving nearly a ten-fold increase over *None*. The resized 3-channel, 30x30 pixel input images equate to 2,700

input values, while the full-sized 3-channel, 224x224 pixel images equate to 150,528 values; this huge reduction in input size results in orders of magnitude fewer tensor operations during CNN inference. Likewise, reducing the color depth of an image from three channels to one reduces the CNN’s computational requirement by two thirds. These transforms, especially resolution reduction, are critical to high query throughput; they enable much smaller CNNs, more than paying for the transforms’ computational costs.

F. Analysis of Increased Cascade Depth

Each additional level added to the cascade exponentially increases the size of the cascade design space. Evaluating the 1.3 million cascades used in results reported elsewhere in this section is fast, taking about 1 minute on average per binary predicate. However, this cascade set includes a full cross-product only for one- and two-level cascades, with three-level cascades restricted to those with fine-tuned ResNet50 as the final classifier. If we considered all possible three-level cascades from all available models, our cascade set balloons to about 45 million distinct cascades, requiring about 40 minutes of evaluation time per predicate. Evaluating a full cross-product of four-level cascades is intractable (roughly 360^4 total cascades).

Figure 11 shows how a cascade set’s Pareto frontier evolves as the maximum depths of the cascades increase. Each set of cascades includes all depths up to its maximum. A “one level” cascade simply corresponds to our set of basic classification models. Where a cascade depth indicates “+ ResNet50”, we append ResNet50 as an additional final level of the cascade. That is, “Two level + ResNet” cascades comprise two levels populated by any model from our collection, followed by a final ResNet50 level. The addition of a ResNet50 level effectively doubles the number of cascades in a cascade set. Adding a full additional level (drawn from our full set of models) increases the cascade set size and evaluation time exponentially.

As can be seen, increasing the depth of the cascades has diminishing returns to throughput at query time, while greatly increasing computational cost of cascade evaluation during system initialization. Moving from “Two level + ResNet50” to “Three level (full)” only increases average throughput by 1.0%, while increasing evaluation time nearly 40-fold to just over 40 minutes. Thus, for our experiments, we have restricted cascades to at most “Two level + ResNet50”. The minimal increase in throughput capabilities of additional layers is not worth the huge increase in evaluation time.

VIII. RELATED WORK

TAHOMA builds upon a rich history of work in the database, image processing, and machine learning communities.

Query systems for visual data — A number of database and query systems targeted at visual data have been developed over the past two decades [21], [42]. Early approaches involved manual textual labeling (e.g. [41]) or extracting rudimentary low-level features (e.g. [13], [39]). Later systems performed additional semantic object extraction, using hand-written functions and statistical methods to define objects [35], [36]. That work, however, precedes the recent deep learning revolution in computer vision and relies on highly manual object extraction methods that limit system applicability and scale.

Work that extended relational query methods over visual data [32], [31], [14] may be useful in extending query capabilities in a system like ours. Fagin’s work with the Garlic system [11], for example, deals with a number of issues particular to querying multimedia databases and the fuzziness of data extracted from multimedia data. However, the data handled by TAHOMA does not exhibit the same kind of fuzziness as Garlic, which is concerned with approximate matches and image similarity. Our system deals with binary predicates that—while based on probabilistic classifier output—are considered strictly true or false at query time.

There has been a recent explosion of interest in query systems for visual data due to advances in classifier accuracy made possible by convolutional neural networks and to the massive datasets made possible by cheap storage and image sensors. NoScope [26], for example, uses several techniques, including classifier cascades, to accelerate query processing over video. Our key departure from this prior work is to include transformations of the input image representation within our query plans, drastically expanding the design space of classifier cascades and enabling much smaller models and order-of-magnitude throughput improvements. We compare our system with NoScope in Section VII-C.

BlazeIt [25] optimizes queries for objects found in video, in part by using small, specialized CNNs to quickly answer queries where possible. Like NoScope, these specialized CNNs do not make use of Tahoma’s input transformations. The rich space of Tahoma’s specialized CNNs, however, could potentially be integrated into a query optimizer like BlazeIt. Focus [19] is a system that indexes objects in live video. It too, uses specialized CNNs to speed up queries. Focus varies the resolution of input images when creating the set of specialized CNNs,

though it does not use TAHOMA’s other input transformations. Additionally, TAHOMA’s use of cascades of multiple specialized CNNs (rather than a single CNN) creates a design space of millions of possible specialized cascades and a much richer set of Pareto-optimal choices. Further, TAHOMA can quickly evaluate its cascades in deployment-specific settings, determining the cascades optimal for the deployment’s current operating characteristics. This could be particularly useful in dynamic scenarios, such as with networked cameras with varying bandwidth constraints. VideoStorm [46] is recent video analytics system that automatically adjusts parameters like video resolution and frame rate to maximize output quality of computer vision algorithms, based on computing resource availability in large clusters.

Image classification — Deep CNNs have revolutionized the field of computer vision, leading to breakthroughs in image classification and detection capabilities in recent years. The ImageNet competition [40] has had as one of its core challenges a 1000-category, million image classification task. A deep CNN was first used in 2012 and greatly reduced the best error rate [17], and the error rate has dropped in the subsequent years to near or better than human performance on the same task (e.g. [23], [16], [20]). Deep CNNs can now facilitate semantic content extraction from images and videos, supporting the development of large scale visual analytics databases. Deep CNNs are too slow, however, in their current incarnations to be applied at scale. Research in speeding up CNNs has seen some recent interest (e.g., SqueezeNet [22] and MobileNets [18]). TAHOMA is essentially classifier-agnostic, so these and future networks can be incorporated into our cascading techniques.

Classifier cascades — Our work leans heavily on previous research into classifier cascades. One of the first works to use the technique was the Viola-Jones face detector [45], which used a series of classifiers based on simple image features to detect faces in photographs. More recently, cascades have been used to accelerate the slow inference speeds of deep neural networks [26], [6], [43]. Other work has used cascades to improve accuracy, rather than speed [44]. Chen et al. [7] used classifier cascades to reduce the cost of expensive feature extraction for text data. Our system, rather, uses feature extraction (in the form of input transformations) to accelerate cascades by creating smaller, faster models. We show how classifier cascade-driven query optimization can be exploited in a visual data analytics system.

Model selection and management — TAHOMA’s general method of generating a large number of model (and cascade) variations and selecting among those can be seen as a form of model selection [29]. A number of recent data systems have been proposed or developed to assist in machine learning model creation, management, and deployment [12], [9], [5], [27]. These systems have been developed with general machine learning tasks in mind, while TAHOMA focuses directly on performing queries and analytics over visual data. Because of this tight focus, we can take advantage of characteristics of CNN architectures and properties of image and video

data. Our model generation and selection methods could be integrated with systems like Velox [9] or MacroBase [5], whose functionality in serving models and handling large amounts of fast data would be complementary to our work.

IX. CONCLUSION AND FUTURE WORK

In this paper, we have presented a method of accelerating content extraction from large corpora of visual data, with the aim of supporting visual analytics queries. We showed how constructing a huge number of classifier cascades from a wide variety of CNN-based classification models can yield large speedups in content extraction. Our cascades applied input transformations on the raw image corpus to further reduce classification costs. We also demonstrated the necessity of being deployment scenario-aware—that is, considering costs such as those for data loading and image transformation—when evaluating the accuracy and throughput tradeoffs of classifiers.

While this paper primarily focused on image classification tasks, it is just the beginning of series of work that will develop a data analytics for visual data that will take full advantage of spatio-temporal locality present in adjacent video frames to further accelerate content extraction. We hope to include new state-of-the-art computer vision methods to extract more complex data, which will then allow the processing of complex analytical queries over video.

REFERENCES

- [1] A year without a byte. <https://code.flickr.net/2017/01/05/a-year-without-a-byte/>. Accessed: 2018-09-04.
- [2] M. Abadi, A. Agarwal, P. Barham, E. Brevdo, Z. Chen, et al. TensorFlow: Large-scale machine learning on heterogeneous systems, 2015. Software available from tensorflow.org.
- [3] S. Agarwal, B. Mozafari, A. Panda, H. Milner, S. Madden, and I. Stoica. BlinkDB: Queries with bounded errors and bounded response times on very large data. In *EuroSys*, 2013.
- [4] M. R. Anderson and M. Cafarella. Input selection for fast feature engineering. In *ICDE*, 2016.
- [5] P. Bailis, E. Gan, S. Madden, D. Narayanan, K. Rong, and S. Suri. Macrobase: Prioritizing attention in fast data. In *SIGMOD*, 2017.
- [6] Z. Cai, M. Saberian, and N. Vasconcelos. Learning complexity-aware cascades for deep pedestrian detection. In *ICCV*, 2015.
- [7] M. Chen, Z. Xu, K. Weinberger, O. Chapelle, and D. Kedem. Classifier cascade for minimizing feature evaluation cost. In *Artificial Intelligence and Statistics*, pages 218–226, 2012.
- [8] F. Chollet et al. Keras. <https://github.com/fchollet/keras>, 2015.
- [9] D. Crankshaw, P. Bailis, J. E. Gonzalez, H. Li, Z. Zhang, M. J. Franklin, A. Ghodsi, and M. I. Jordan. The missing piece in complex analytics: Low latency, scalable model management and serving with velox. *arXiv preprint arXiv:1409.3809*, 2014.
- [10] Facebook, Ericsson, and Qualcomm. A Focus on Efficiency (whitepaper). Technical report, 2013.
- [11] R. Fagin. Fuzzy queries in multimedia database systems. In *PODS*, 1998.
- [12] X. Feng, A. Kumar, B. Recht, and C. Ré. Towards a unified architecture for in-RDBMS analytics. In *SIGMOD*, 2012.
- [13] M. Flickner, H. Sawhney, W. Niblack, et al. Query by image and video content: The QBIC system. *Computer*, 28(9):23–32, 1995.
- [14] F. Golshani and N. Dimitrova. A language for content-based video retrieval. *Multimedia Tools and Applications*, 6(3):289–312, 1998.
- [15] I. Goodfellow, Y. Bengio, and A. Courville. *Deep Learning*. MIT Press, 2016. <http://www.deeplearningbook.org>.
- [16] K. He, X. Zhang, S. Ren, and J. Sun. Deep residual learning for image recognition. In *CVPR*, 2016.
- [17] G. Hinton, O. Vinyals, and J. Dean. Distilling the knowledge in a neural network. In *NIPS Deep Learning Workshop*, 2014.
- [18] A. G. Howard, M. Zhu, B. Chen, D. Kalenichenko, W. Wang, T. Weyand, M. Andreetto, and H. Adam. Mobilenets: Efficient convolutional neural networks for mobile vision applications. *arXiv:1704.04861*, 2017.
- [19] K. Hsieh, G. Ananthanarayanan, P. Bodik, P. Bahl, M. Philipose, P. B. Gibbons, and O. Mutlu. Focus: Querying large video datasets with low latency and low cost. In *OSDI*, 2018.
- [20] J. Hu, L. Shen, and G. Sun. Squeeze-and-excitation networks. *arXiv preprint arXiv:1709.01507*, 2017.
- [21] W. Hu, N. Xie, L. Li, X. Zeng, and S. Maybank. A survey on visual content-based video indexing and retrieval. *IEEE Trans. on Systems, Man, and Cybernetics, Part C*, 41(6):797–819, 2011.
- [22] F. N. Iandola, S. Han, M. W. Moskewicz, K. Ashraf, W. J. Dally, and K. Keutzer. SqueezeNet: AlexNet-level accuracy with 50x fewer parameters and <0.5mb model size. *arXiv:1602.07360*, 2016.
- [23] S. Ioffe and C. Szegedy. Batch normalization: Accelerating deep network training by reducing internal covariate shift. In *ICML*, 2015.
- [24] Y. Jia, E. Shelhamer, J. Donahue, S. Karayev, J. Long, R. Girshick, S. Guadarrama, and T. Darrell. Caffe: Convolutional architecture for fast feature embedding. In *Intl. Conference on Multimedia*. ACM, 2014.
- [25] D. Kang, P. Bailis, and M. Zaharia. Blazelt: Fast exploratory video queries using neural networks. *arXiv preprint arXiv:1805.01046*, 2018.
- [26] D. Kang, J. Emmons, F. Abuzaid, P. Bailis, and M. Zaharia. NoScope: Optimizing neural network queries over video at scale. *Proceedings of the VLDB Endowment*, 10(11):1586–1597, 2017.
- [27] T. Kraska, A. Talwalkar, J. C. Duchi, R. Griffith, M. J. Franklin, and M. I. Jordan. MLBase: A distributed machine-learning system. In *CIDR*, 2013.
- [28] A. Krizhevsky, I. Sutskever, and G. E. Hinton. Imagenet classification with deep convolutional neural networks. In *NIPS*, 2012.
- [29] A. Kumar, R. McCann, J. Naughton, and J. M. Patel. Model selection management systems: The next frontier of advanced analytics. *ACM SIGMOD Record*, 44(4):17–22, 2016.
- [30] H.-T. Kung, F. Luccio, and F. P. Preparata. On finding the maxima of a set of vectors. *Journal of the ACM (JACM)*, 22(4):469–476, 1975.
- [31] T. C. Kuo and A. L. Chen. A content-based query language for video databases. In *Multimedia Computing and Systems*, 1996.
- [32] J. Z. Li, M. T. Ozsu, D. Szafron, and V. Oria. MOQL: A multimedia object query language. In *Intl. Workshop on Multimedia Info. Sys.*, 1997.
- [33] C. H. Papadimitriou and M. Yannakakis. Multiobjective query optimization. In *PODS*, 2001.
- [34] Y. Park, B. Mozafari, J. Sorenson, and J. Wang. VerdictDB: Universalizing approximate query processing. In *SIGMOD*, 2018.
- [35] M. Petkovic and W. Jonker. *Content-Based Video Retrieval: A Database Perspective*, volume 11. Springer, 1999.
- [36] M. Petkovic and W. Jonker. A framework for video modelling. In *Intl. Conf. on Applied Informatics*, 2000.
- [37] J. Redmon, S. Divvala, R. Girshick, and A. Farhadi. You only look once: Unified, real-time object detection. In *CVPR*, 2016.
- [38] J. Redmon and A. Farhadi. YOLO9000: Better, faster, stronger. In *CVPR*, 2017.
- [39] W. Ren, S. Singh, M. Singh, and Y. S. Zhu. State-of-the-art on spatio-temporal information-based video retrieval. *Pattern Recognition*, 42(2):267–282, 2009.
- [40] O. Russakovsky, J. Deng, H. Su, J. Krause, S. Satheesh, S. Ma, Z. Huang, A. Karpathy, A. Khosla, M. Bernstein, A. C. Berg, and L. Fei-Fei. ImageNet Large Scale Visual Recognition Challenge. *International Journal of Computer Vision (IJCV)*, 115(3):211–252, 2015.
- [41] T. G. A. Smith and G. Davenport. The stratification system a design environment for random access video. In *Intl. Workshop on Network and Operating System Support for Digital Audio and Video*, 1992.
- [42] C. G. Snoek and M. Worring. Concept-based video retrieval. *Foundations and Trends in Information Retrieval*, 2(4):215–322, 2008.
- [43] Y. Sun, X. Wang, and X. Tang. Deep convolutional network cascade for facial point detection. In *CVPR*, 2013.
- [44] R. Verschae, J. Ruiz-del Solar, and M. Correa. A unified learning framework for object detection and classification using nested cascades of boosted classifiers. *Machine Vision and Apps.*, 19(2):85–103, 2008.
- [45] P. Viola and M. J. Jones. Robust real-time face detection. *International Journal of Computer Vision*, 57(2):137–154, 2004.
- [46] H. Zhang, G. Ananthanarayanan, P. Bodik, M. Philipose, P. Bahl, and M. J. Freedman. Live video analytics at scale with approximation and delay-tolerance. In *NSDI*, 2017.