

Name Your Colour For the Task: Artificially Discover Colour Naming via Colour Quantisation Transformer

Shenghan Su¹, Lin Gu^{3,2*}, Yue Yang^{1,4}, Zenghui Zhang¹, Tatsuya Harada^{2,3}

¹Shanghai Jiao Tong University, ²The University of Tokyo, ³RIKEN AIP, ⁴Shanghai AI Laboratory

{su2564468850, yang-yue, zenghui.zhang}@sjtu.edu.cn, lin.gu@riken.jp, harada@mi.t.u-tokyo.ac.jp




Figure 1: (a) The theoretical limit of efficiency for colour naming (black curve) and cases of the WCS probability map of human colour language copied from [49]. (b) The colour size (from 1-bit to 6-bit)-accuracy curve on the tiny-imagenet-200 [29] dataset. The WCS probability maps generated by our CQFormer are also shown along the curve. (c) The colour naming stimulus grid used in the WCS [26]. (d) The three-term WCS probability map of CQFormer after embedding 1978 Nafaanra three-colour system ((light ('fiNge'), dark ('wOO'), and warm or red-like ('nyiE'))) into the latent representation. (e) The four-term WCS probability map of CQFormer evolved from (d). The evolved fourth colour, yellow-green, is consistent with the prediction of basic colour term theory [2].

Abstract

The long-standing theory that a colour-naming system evolves under dual pressure of efficient communication and perceptual mechanism is supported by more and more linguistic studies, including analysing four decades of diachronic data from the Nafaanra language. This inspires us to explore whether machine learning could evolve and discover a similar colour-naming system via optimising the communication efficiency represented by high-level recognition performance. Here, we propose a novel colour quantisation transformer, CQFormer, that quantises colour space while maintaining the accuracy of machine recognition on the quantised images. Given an RGB image, Annotation Branch maps it into an index map before generating the quantised image with a colour palette; meanwhile the Palette Branch utilises a key-point detection way to find proper colours in the palette among the whole colour space. By interacting with colour annotation, CQFormer is able to balance both the machine vision accuracy and colour perceptual structure

such as distinct and stable colour distribution for discovered colour system. Very interestingly, we even observe the consistent evolution pattern between our artificial colour system and basic colour terms across human languages. Besides, our colour quantisation method also offers an efficient quantisation method that effectively compresses the image storage while maintaining high performance in high-level recognition tasks such as classification and detection. Extensive experiments demonstrate the superior performance of our method with extremely low bit-rate colours, showing potential to integrate into quantisation network to quantities from image to network activation. The source code is available at <https://github.com/ryeochiv/CQFormer>

1. Introduction

Hath not a Jew eyes? Hath not a Jew hands, organs,
dimensions, senses, affections, passions?

William Shakespeare "The Merchant of Venice"

Does artificial intelligence share the same perceptual mechanism for colours as human beings? We aim to explore

*Corresponding author.

this intriguing problem from machine learning perspective.

Colour involves the visual reception and neural registering of light stimulants when the spectrum of light interacts with cone cells in the eyes. Physical specifications of colour also include the reflective properties of the physical objects, geometry incident illumination, *etc.* By defining a colour space [18], people could identify colours directly according to these quantifiable coordinates.

Compared to the pure physiological nature of hue categorisation, the complex phenomenon of colour naming or colour categorisation spans multiple disciplines, from cognitive science to anthropology. Solid diachronic research [2] also suggests that human languages constantly evolve to acquire new colour names, resulting in an increasingly fine-grained colour naming system. This evolutionary process is hypothesised to be under the pressure of communication efficiency and perceptual structure. Communication efficiency requires shared colour partitioning to be communicated accurately with a lexicon as simple and economical as possible. Colour perceptual structure is relevant to human perception in nature. For example, the colour space distance between nearby colours should correspond to their perceptual dissimilarity. This structure of perceptual colour space has long been used to explain colour naming patterns across languages. A recent analysis of human colour naming systems, especially in Nafaanra, contributes very direct evidence to support this hypothesis through the employment of Shannon’s communication model [41]. Interestingly, this echoes the research on colour quantisation, which quantises colour space to reduce the number of distinct colours in an image.

Traditional colour quantisation methods [23, 19, 17] are *perception-centred* and generate a new image that is as visually perceptually similar as possible to the original image. These methods group similar colours in the colour space and represent each group with a new colour, thus naturally preserving the perceptual structure. Instead of prioritising the perceptual quality, Hou *et al.* [25] proposed a *task-centred/machine-centred* colour quantisation method, ColorCNN, focusing on maintaining network classification accuracy in the restricted colour spaces. While achieving an impressive classification accuracy on even a few-bit image, ColorCNN only identifies and preserves *machine-centred* structure without directly clustering similar colours in the colour space. Therefore, this pure *machine-centred* strategy sacrifices perceptual structure and often associates nearby colours with different quantised indices.

Zaslavsky *et al.* [49] measure the communication efficiency in colour naming by analysing the informational complexity based on the information bottleneck (IB) principle. Here, we argue that the network recognition accuracy also reflects the communication efficiency when the number of colours is restricted. Since the human colour naming is shaped by both perception structure and communication effi-

ciency [51], we integrate the need for both *perception* and *machine* to propose a novel end-to-end colour quantisation transformer, CQFormer, to discover the artificial colour naming systems.

As illustrated in Fig. 1 (b), the recognition accuracy increases with the number of colours in our discovered colour naming system. Surprisingly, the complexity-accuracy trade-offs are similar to the numerical results (Fig. 1 (a)) independently derived from linguistic research [49]. What is more, after embedding 1978 Nafaanra three-colour system (Nafaanra-1978) into the latent representation of CQFormer (Fig. 1 (d)), our method automatically evolves the fourth colour closed to yellow-green, matching the basic colour terms theory [2] summarised in different languages. Berlin and Kay found universal restrictions on colour naming across cultures and claimed languages acquire new basic colour categories in a strict chronological sequence. For example, if a culture has three colours (light (‘fiNge’), dark (‘wOO’), and warm or red-like (‘nyiE’) in Nafaanra), the fourth colour it evolves should be yellow or green, exactly the one (Fig. 1 (e)) discovered by our CQFormer.

The pipeline of CQFormer, shown in Fig. 2, comprises two branches: Annotation Branch and Palette Branch. Annotation Branch annotates each pixel of the input RGB image with the proper quantised colour index before painting the index map with the corresponding colour in the colour palette. We localise the colour palette in the whole RGB colour space with a novel Palette Branch, which detects the key-point with explicit attention queries of transformer. During the training stage, as illustrated in the red line and black line of Fig. 2 (a), Palette Branch interacts with an input image and Reference Palette Queries to maintain the perceptual structure by reducing the perceptual structure loss. This *perception-centred* design groups similar colours and ensures the colour palette sufficiently represents the colour naming system defined by the World Color Survey (WCS) colour naming stimulus grids. As shown in Fig. 2 (b), each item in the colour palette (noted by an asteroid) lies in the middle of the corresponding colour distribution in the WCS colour naming probability map. Finally, the quantised image is passed to a high-level recognition module for *machine accuracy* tasks such as classification and detection. Through the joint optimisation of CQFormer and consequent high-level module, we can balance both *perception* and *machine*. Besides automatically discovering the colour naming system, our CQFormer also offers an effective solution to extremely compress image storage while maintaining high performance in high-level recognition tasks. For example, the CQFormer achieves 50.6% top-1 accuracy on the CIFAR100 [28] dataset with only a 1-bit colour space (*i.e.*, two colours). The extremely low-bit quantisation of our also demonstrates the potential to integrate into quantisation network research [40, 47], allowing the end-to-end optimisation from image to weight and

activation.

Our contributions could be summarised as follows:

- We propose a novel end-to-end colour quantisation transformer, CQFormer, to artificially discover a colour naming system by counting both perception and machine. The discovered colour naming system shows a similar pattern to human language on colour.
- We propose a novel colour quantisation method taking palette generation as an attention-based key-point detection task. With the input of independent attention queries, it automatically generates 3D coordinates as the selected colour in the whole colour space.
- Our colour quantisation achieves superior classification and object detection performance with extremely low bit-rate colours.

2. Related Works

Colour Quantisation. Colour quantisation, also known as optimal palette generation, compresses images by remapping original pixel colours to a limited set of entries in a small palette.

Traditional colour quantisation [37, 14, 32, 46] aims to reduce the colour space while preserving visual fidelity. These *perception-centred* methods usually cluster colours to create a new image as visually perceptually similar as possible to the original image. For example, MedianCut [23] and OC-Tree [19] are two representative clustering-based algorithms. Dithering [16] eliminates visual artefacts by including a noise pattern. The colour-quantised images can be expressed as indexed colour [39], and encoded with PNG [3].

There are also much of efforts [7, 8] to generate the optimal colour palette and recolour the image. For example, Bahng *et al.* [1] proposed the text-to-palette generation networks to generate an appropriate palette according to the semantics of the text input. Yoo *et al.* [48] leverage the memory networks to retrieve the feature of colour palette for colourisation with limited data. Li *et al.* [30] develop a self-supervised approach to recolouring images from design-oriented fields by reproducing colour palettes with luminance distributions different from the input.

Recently, Hou *et al.* [25] propose a pure *machine-centred* CNN-based colour quantisation method, ColorCNN, which effectively maintains the informative structures under an extremely constrained colour space. In addition to colour quantisation, Camposeco *et al.* [4] also design a task-centred image compression method for localisation in 3D map data. Since human colour naming reflects both perceptual structure and communicative need, our CQFormer also considers both *perception* and *machine* to discover the colour naming system artificially.

World Color Survey. The World Color Survey (WCS) [26] comprises colour name data from 110 languages of non-industrialised societies [52], *w.r.t.* the stimulus grid shown in Fig. 1 (c). There are 320 colour chips in colour naming stimulus grids, and each chip is at its maximum saturation for that hue/lightness combination, while columns correspond to equally spaced Munsell hues and rows to equally spaced Munsell values. Participants were asked to name the colour of each chip to record the colour naming system, generating the WCS probability map for human language (*e.g.* the Nafaanra-1978 in Fig. 3 (c)).

Colour Categorisation/Naming. Van De Weijer *et al.* [45] gathered datasets from Google and Ebay with the explicit aim of learning colour names from real-world images. Paraga and Akbarinia[38] then recruited 17 paid subjects and employed a psychophysical experiment to bridge the gap between the physiology of the visual system and colour categorisation. Gibson *et al.* [20] indicated that warm colours are communicated more efficiently than cool colours in general, and crucially, categorical variations between languages are attributed to differences in the practical value of colour. Siuda-Krzywicka *et al.*[42] indicates the independence of colour categorisation and naming, while also offering a plausible neural foundation for the process of colour naming. Chaabouni *et al.*[6] treated communicative concepts as the primary driving force behind the formation of colour categories and demonstrated that continuous message passing increased system complexity and reduced efficiency. De Vries *et al.* [12] demonstrated that an artificial visual system develops colour categorical boundaries resembling those of humans by learning to recognise objects in images.

3. Methodology

3.1. Problem Formulation

For a dataset \mathcal{D} containing image-label pairs (x, y) , the recognition network $f_{\theta}(\cdot)$ takes the input image x and predicts its label $\hat{y} = f_{\theta}(x)$ (*e.g.*, the class probability in a classification task). The network’s parameters θ can be optimised by minimising the loss between the predicted label \hat{y} and the ground truth label y , defined as the machine-centred loss \mathcal{L}_M . This process allows us to find the optimal parameter set θ^* .

$$\theta^* = \arg \min_{\theta} \sum_{(x,y) \in \mathcal{D}} \mathcal{L}_M(y, f_{\theta}(x)). \quad (1)$$

We aim to discover the artificial colour naming system that balances the need for both machine accuracy and human perception. To achieve this, CQFormer focuses not only on recognition accuracy but also on preserving the perceptual structure of the image. The function in Eq. 1 can be rewritten as follows:

$$\psi^*, \theta^* = \arg \min_{\psi, \theta} \sum_{(x,y) \in \mathcal{D}} \mathcal{L}_M(y, f_{\theta}(g_{\psi}(x))) + \mathcal{L}_P, \quad (2)$$

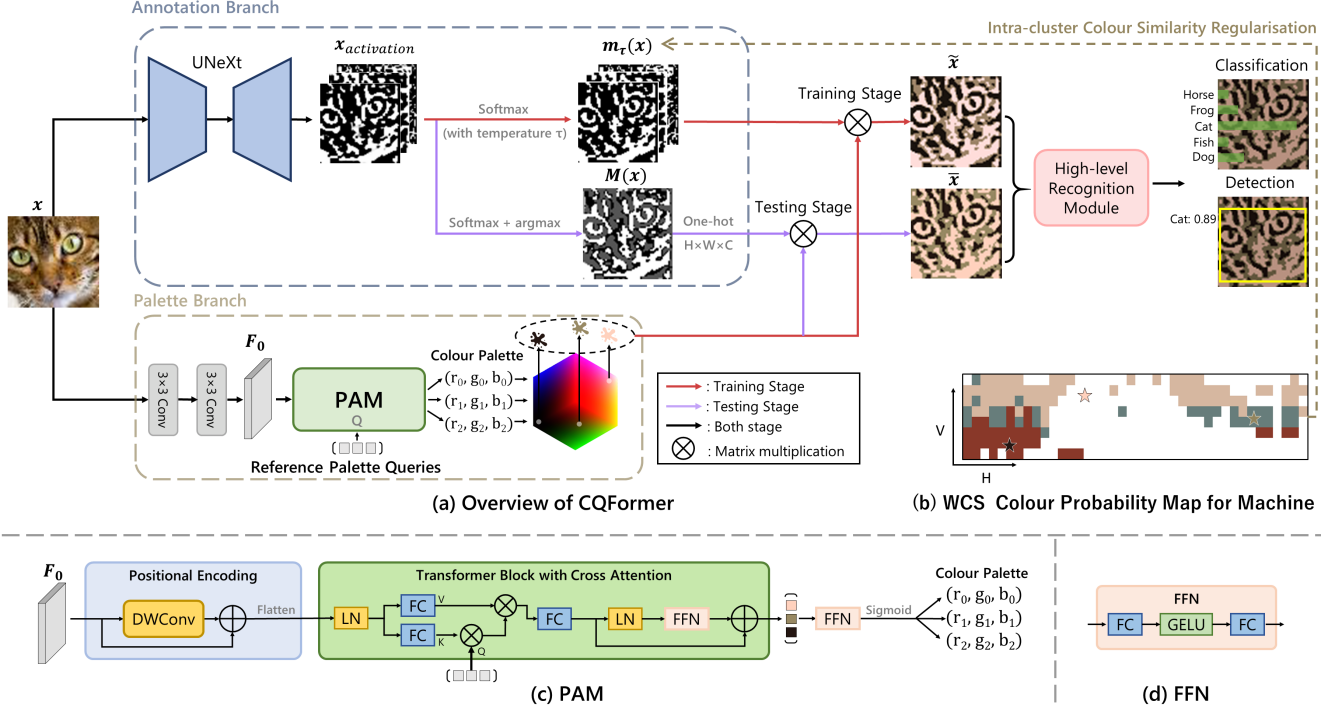


Figure 2: (a) Overview of our CQFormer (taking 3-colour quantisation as an example). Red lines are only for training stage, purple lines are only for testing stage, and black lines are both used in both stage. (b) WCS colour probability map counted across all pixels of the input image. (c) Detailed structure of PAM. (d) Detailed structure of FFN.

where ψ and θ represent the parameters of the CQFormer g and the high-level recognition network f , respectively. We jointly optimise g and f to find the optimal parameters ψ^* and θ^* . Additionally, \mathcal{L}_P is a perceptual structure loss that is focused on preserving the perceptual structure of the image and will be further explained in detail in Sec. 3.4.

3.2. CQFormer Architecture

Overall Architecture. An overview of the CQFormer is illustrated in Fig. 2 (a). It consists of two main branches: (1) Annotation Branch, which assigns a quantised colour index to each pixel of the input RGB image, and (2) Palette Branch, which is responsible for generating a suitable colour palette.

Annotation Branch of CQFormer takes an original image $\mathbf{x} \in \mathbb{R}^{H \times W \times 3}$ as input, where H and W are the height and width of the image, respectively. During the training stage, it generates a probability map $m_\tau(\mathbf{x}) \in \mathbb{R}^{H \times W \times C}$, where C is the number of quantised colours, and $\tau > 0$ is the temperature parameter of the Softmax function [22]. During the testing stage, it produces a one-hot colour index map $\text{One-hot}(M(\mathbf{x})) \in \mathbb{R}^{H \times W \times C}$, where each pixel of the image is assigned a single colour index among the C quantised colours.

Palette Branch of CQFormer takes the original image \mathbf{x} and Reference Palette Queries $\mathbf{Q} \in \mathbb{R}^{C \times d}$ as input. These queries are composed of C learnable vectors of dimension d , each representing an automatically mined colour. The

queries \mathbf{Q} interact with the keys $\mathbf{K} \in \mathbb{R}^{(\frac{HW}{16}) \times d}$ and values $\mathbf{V} \in \mathbb{R}^{(\frac{HW}{16}) \times d}$, generated from the input image \mathbf{x} to produce the colour palette $P(\mathbf{x}) \in \mathbb{R}^{C \times 3}$. This palette consists of C triples (R, G, B) , each representing one of the machine-discovered C colours.

Finally, CQFormer produces the quantised image by performing a matrix multiplication between $m_\tau(\mathbf{x})$ and $P(\mathbf{x})$ during the training stage. During the testing stage, it is obtained from $\text{One-hot}(M(\mathbf{x}))$ and $P(\mathbf{x})$. The quantised image is then fed into the high-level recognition module for object detection and classification.

Annotation Branch. The first component of Annotation Branch is a UNeXt encoder [44] that generates per-pixel categories. Given the input image \mathbf{x} , the encoder produces a class activation map $\mathbf{x}_{activation} \in \mathbb{R}^{H \times W \times C}$, which contains crucial and semantically rich features. Then, different approaches are used to process the class activation map during testing and training.

(1) *Testing stage.* As shown by the purple lines in Fig. 2 (a), we use the $\mathbf{x}_{activation}$ as the input to a Softmax function over C channels, and then apply an arg max function to produce a colour index map $M(\mathbf{x}) \in \mathbb{R}^{H \times W}$:

$$M(\mathbf{x}) = \arg \max_C (\text{Softmax}(\mathbf{x}_{activation})). \quad (3)$$

Subsequently, the colour index map is transformed into a one-hot encoding, denoted as $\text{One-hot}(M(\mathbf{x})) \in \mathbb{R}^{H \times W \times C}$, which is then combined with the colour palette

$P(\mathbf{x})$ through matrix multiplication, resulting in the generation of the test-time colour-quantised image $\tilde{\mathbf{x}}$.

$$\tilde{\mathbf{x}} = \text{One-hot}(M(\mathbf{x})) \otimes P(\mathbf{x}), \quad (4)$$

where \otimes represents matrix multiplication.

(2) *Training stage.* As depicted by the red lines in Fig.2 (a), since the $\arg \max$ function is not differentiable, we use the Softmax function as a substitute during the training stage. To prevent overfitting, we incorporate a temperature parameter τ [22] into the Softmax function, pushing the probability map distribution closer to a one-hot vector. This results in the probability map $m_\tau(\mathbf{x})$ with temperature parameter τ , computed as:

$$m_\tau(\mathbf{x}) = \text{Softmax}\left(\frac{\mathbf{x}_{\text{activation}}}{\tau}\right). \quad (5)$$

As extensively discussed in [22], the output is similar to a one-hot vector with large diversity when the temperature is low ($0 < \tau < 1$) and a uniform distribution with small diversity otherwise ($\tau > 1$). Therefore, we set $0 < \tau < 1$ and the train-time colour-quantised image $\tilde{\mathbf{x}}$ is generated as:

$$\tilde{\mathbf{x}} = m_\tau(\mathbf{x}) \otimes P(\mathbf{x}). \quad (6)$$

Palette Branch. We locate the representative colours by using an attention-based key-point detection strategy, which is originally designed to use transformer queries to automatically find the key-point location, such as bounding-boxes [5], human poses [31] and colour matrix with gamma [11]. In other words, we reformulate the problem of colour quantisation into a 3D spatial key-point localisation task within the whole RGB colour space.

Given the input image \mathbf{x} , we first extract a high-dimensional lower resolution feature $F_0 \in \mathbb{R}^{\frac{H}{4} \times \frac{W}{4} \times d}$ using two stacked convolution layers. The F_0 is then passed to the Palette Acquisition Module (PAM) to acquire the colour palette $P(\mathbf{x})$. As shown in Fig. 2 (c), different from DETR [5], our $\mathbf{Q} \in \mathbb{R}^{C \times d}$, called Reference Palette Queries, are explicit learnable embeddings without extra multi-head self-attention, which attends \mathbf{K} and \mathbf{V} generated from F_0 . The positional encoding mechanism is a depth-wise convolution [9], which is suitable for different input resolutions. After that, the position-encoded feature is flattened into sequences before being passed into our transformer block. Here, the \mathbf{K} and \mathbf{V} are generated by two fully-connection (FC) layers separately, and the cross-attention is estimated as:

$$\text{Attention}(\mathbf{Q}, \mathbf{K}, \mathbf{V}) = \text{Softmax}\left(\frac{\mathbf{Q}\mathbf{K}^\top}{\sqrt{d}}\right)\mathbf{V}. \quad (7)$$

After a feed forward network (FFN) [15], including two FC layers and a GELU [24] activation function, then with a residual connection [21], the Reference Palette Queries \mathbf{Q} are transformed into output embeddings. We then decode embeddings into 3D coordinates with another FFN and a

Sigmoid function, resulting in the final localisation of colour palette $P(\mathbf{x}) \in \mathbb{R}^{C \times 3}$ within the whole RGB colour space.

High-Level Recognition Module. Taking the quantised image ($\tilde{\mathbf{x}}$ or $\tilde{\tilde{\mathbf{x}}}$) as input, the high-level recognition module predicts the results of object detection or classification.

3.3. WCS Colour Probability Map

As introduced in WCS of Sec. 2, the WCS probability map $m_{\text{WCS}} \in \mathbb{R}^{8 \times 40 \times C}$ for human language is collected by the participants' colour perception in WCS. We also define the WCS colour probability map for machine (see in Fig. 1 (d) (e)) and generate it from the datasets. First, the colour index of each pixel is mapped to the grid with the closest (hue, value) coordinate as the pixel. Second, we count the frequency of occurrence of each colour index in each grid, and select the pixel's colour index with the highest frequency as the grid's colour index. Therefore, C clusters are formed according to the C grid's colour indexes in WCS colour probability map. Finally, each cluster is shown in the colour corresponding to the centre of mass of its colour category.

The WCS colour probability map for machine is then used to measure the similarity of colours within the same cluster in Sec.3.4. More importantly, it creates a strong correlation between human language and machine vision since they can be represented in the same format. Therefore, we explore the colour evolution using the WCS colour probability map in Sec.3.5.

3.4. Perceptual Structure Loss

To ensure that CQFormer also produces visually pleasing and structurally sound results, an additional perceptual structure loss \mathcal{L}_P is added to the training process:

$$\mathcal{L}_P = \alpha R_{\text{Colour}} + \beta R_{\text{Diversity}} + \gamma \mathcal{L}_{\text{Perceptual}}. \quad (8)$$

α , β and γ are three non-negative parameters, which combine intra-cluster colour similarity regularisation R_{Colour} , diversity regularisation $R_{\text{Diversity}}$ and perceptual similarity loss $\mathcal{L}_{\text{Perceptual}}$, respectively. Therefore, the objective of training the CQFormer is to minimise the total loss $\mathcal{L}_{\text{total}}$:

$$\mathcal{L}_{\text{total}} = \mathcal{L}_M + \mathcal{L}_P. \quad (9)$$

Intra-cluster Colour Similarity Regularisation. The CQFormer associates each pixel of the input image with a colour index and forms C clusters covering different parts of the WCS colour probability map. To ensure that the pixels within the same cluster are perceptually similar in colour, we propose the intra-cluster colour similarity regularisation R_{Colour} , which is the mean of the colour variance within the C clusters. At first, the centroid colour values of the C clusters $\{\mu_1, \mu_2, \dots, \mu_C\}$ are calculated. Then, we compute C squared colour distances d_{HSV}^2 between all pixels and their centroid colour value in each cluster. Here, the d_{HSV}^2 is

calculated in Munsell HSV colour space as:

$$d_{\text{HSV}}^2(h_1, s_1, v_1, h_2, s_2, v_2) = (v_2 - v_1)^2 + s_1^2 v_1^2 + s_2^2 v_2^2 - 2s_1 s_2 v_1 v_2 \cos(h_2 - h_1). \quad (10)$$

Finally, we take the mean value of the C squared distances as the R_{Colour} :

$$R_{\text{Colour}} = \frac{1}{C} \times \sum_{c=1}^C \frac{1}{N_c} \sum_{i=1}^{N_c} d_{\text{HSV}}^2(\mathbf{x}_c[i], \mu_c), \quad (11)$$

where N_c is the number of all pixels in $Cluster_c$, and $\mathbf{x}_c[i]$ represents the Munsell HSV value of the i -th pixel.

Diversity Regularisation. To encourage the CQFormer to select at least one pixel of all C colours, we adopt the diversity regularisation term $R_{\text{Diversity}}$ proposed by Hou *et al.* [25]. Diversity is a simple yet efficient metric and serves as an implicit entropy of the assignment distribution. The $R_{\text{Diversity}}$ is calculated as:

$$R_{\text{Diversity}} = \log_2 C \times \left(1 - \frac{1}{C} \times \sum_c \max_{(u,v)} [m_\tau(\mathbf{x})]_{u,v} \right). \quad (12)$$

Perceptual Similarity Loss. The perceptual similarity loss, denoted by $\mathcal{L}_{\text{Perceptual}}$, is a mean squared error (MSE) loss between the quantised image $\tilde{\mathbf{x}}$ and the input image \mathbf{x} . It ensures that each item of the colour palette $P(\mathbf{x})$ lies in the centre of corresponding colour distribution in the WCS colour probability map (noted by asteroids in Fig. 2 (b)). The $\mathcal{L}_{\text{Perceptual}}$ is calculated as:

$$\mathcal{L}_{\text{Perceptual}} = \mathcal{L}_{\text{MSE}}(\tilde{\mathbf{x}}, \mathbf{x}). \quad (13)$$

3.5. Colour Evolution

With CQFormer, we explore the colour evolution based on the classification task, involving two successive stages with different loss functions. Since the CQFormer initially has no prior knowledge of colour naming systems associated with corresponding human languages, the first embedding stage aims to embed the colour perceptual knowledge of a certain language into the latent representation of the CQFormer. For example, the CQFormer first learns and matches the 1978 Nafaanra three-colour system by forcing the CQFormer to output a similar WCS colour probability map to that for Nafaanra. Here, we design two embedding solutions and loss functions, *i.e.* $\mathcal{L}_{\text{Full-Embedding}}$ and $\mathcal{L}_{\text{Central-Embedding}}$, to distil either full colour probability map embedding or only representative colours to CQFormer. The second evolution stage then lets CQFormer evolve more colours, *i.e.* splitting the fourth colour from the learned three-colour system under the pressure of both accuracy and perceptual structure.

Embedding Stage.

(1) *Colour Probability Map Embedding.* As illustrated in Fig. 3, this embedding forces our CQFormer to match the




Figure 3: The procedure of colour probability map embedding. (a) The origin RGB image. (b) The WCS colour naming stimulus grids. (c) The m_{WCS} of Nafaanra-1978 copied from [49]. (d) The human language probability map after $\arg \max$ function for Nafaanra-1978: $\arg \max(m_{\text{Human}}(\mathbf{x}))$.

identical WCS colour probability map of a certain language. At first, we ensure that the CQFormer outputs the same number of quantised colours as the certain human colour system. For each pixel in the input image, we collect its Munsell (hue, value) coordinate and spatial position coordinate (i, j) in the input image. Then we locate this (hue, value) coordinate on the WCS colour probability map of a certain language (Fig. 3 (a)→(b)→(c)) to find its corresponding C probability values $p_{\text{Human}}(i, j) \in \mathbb{R}^C$ of human colour categories, *e.g.* 9% for ‘fiNge’, 4% for ‘wOO’, and 87% for ‘nyiE’. After performing the above operations on all pixels of the input image, a set of probability values $\{p_{\text{Human}}(i, j)\}_{(i,j)=(1,1)}^{(H,W)}$ are generated. We arrange each of them according to their spatial position coordinates (i, j) to obtain a new matrix regarded as human language probability map $m_{\text{Human}}(\mathbf{x}) \in \mathbb{R}^{H \times W \times C}$. Then we use a cross-entropy loss \mathcal{L}_{CE} instead of \mathcal{L}_P . Finally, for this full embedding solution, we minimise the loss function as:

$$\mathcal{L}_{\text{Full-Embedding}} = \mathcal{L}_M + \mathcal{L}_{\text{CE}}(m_\tau(\mathbf{x}), m_{\text{Human}}(\mathbf{x})). \quad (14)$$

By minimising the above Eq. 14, our CQFormer inherits the full colour naming system of human language.

(2) *Central Colour Embedding.* Alternatively, we could distil less colour naming information. Here, we only embed representative colours and their C central colour (hue, value) coordinates $\{\mu_{\text{Human},c}\}_{c=1}^C$ in the WCS colour probability map of the certain human language (see three asteroids in Fig. 3 (c)). For this embedding solution, we minimise the loss function as:

$$\mathcal{L}_{\text{Central-Embedding}} = \mathcal{L}_M + R_{\text{Colour}}, \quad (15)$$

where we replace the μ_c in R_{Colour} with aforementioned $\mu_{\text{Human},c}$ and ignore the saturation of Munsell HSV by setting $s_1 = s_2 = 1$.

Colour Evolution Stage. After the distillation of a certain colour naming system of human language, we evolve the CQFormer to split fine-grained new colour. In this stage, we remove the restriction on the number of quantised colours and minimise different combinations of losses to encourage the CQFormer to evolve more colours. Please refer to Sec. 4.4 for more details.

Table 1: Object detection results on MS COCO dataset [34] with Sparse-RCNN [43] detector, here we report the average precision (AP) value.

Method	1-bit	2-bit	3-bit	4-bit	5-bit	6-bit	Full Colour (24-bit)
Upper bound	-	-	-	-	-	-	45.0
Median Cut (w/o D) [23]	11.5	12.7	15.4	17.0	20.4	23.2	-
Median Cut (w D) [17]	12.3	13.8	15.2	19.6	21.8	25.6	-
OCTree [19]	10.7	13.4	13.2	16.7	18.9	22.8	-
CQFormer	13.9	16.5	18.8	21.5	27.5	29.8	-

4. Experiments

We evaluate the CQFormer on mainstream benchmark datasets of both object detection task (Sec. 4.2) and image classification task (Sec. 4.3). Additionally, we specifically design a colour evolution experiment (Sec. 4.4) to demonstrate how our CQFormer automatically evolves to increase fine-grained colours. For ablation study, visualisation and other details, please refer to Supplementary Material.

4.1. Datasets and Experiment Settings

Datasets. For object detection, we utilise MS COCO [34] dataset, which contains $\sim 118k$ images with bounding box annotation in 80 categories. Here, we use COCO `train2017` set as the train set and use COCO `val2017` set as the test set. For classification, we use CIFAR10, CIFAR100 [28], STL10 [10] and tiny-imagenet-200 (Tiny200) [29] dataset. Both CIFAR10 and CIFAR100 contain 60000 images, covering 10 and 100 classes of general objects, respectively. The STL10 consists of 13000 images (5000 for training and 8000 for testing) with 10 classes. The Tiny200 is a subset of ImageNet [13] and contains 200 categories with 110k images. We adopt random crop and random flip for data augmentation.

Evaluation Metrics. For object detection, we report average precision (AP) value. For classification, we report top-1 classification accuracy.

Implement Details.

(1) *Upper bound.* We utilise the performance of the detector/classifier without additional colour quantisation methods in full-colour space (24-bit) as the upper bound. For object detection, we adopt the Sparse-RCNN [43] detector with ResNet-50 [21] backbone and FPN [33] neck. For classification, we adopt Resnet-18 [21] network.

(2) *Training Settings.* All colour quantisation experiments are finished at quantisation levels from 1-bit to 6-bit, *i.e.* $C \in \{2, 4, 8, 16, 32, 64\}$. We set the temperature parameter $\tau = 0.01$ and the $\alpha = 1, \beta = 0.3, \gamma = 1$. For detection, only the Resnet-50 backbone of the detector is initialised with Imagenet [13] pre-train weights. We jointly train our CQFormer with Sparse-RCNN for 36 epochs on 4 Tesla V100 GPUs. The batch size is set to 8, and the optimiser is AdamW [36]. The initial learning rate is set to $1.25e^{-6}$,

and the learning rate would decay to one-tenth at 24 epochs. For classification, we cascade CQFormer with the Resnet-18 and jointly train them without pre-trained models on a single GeForce RTX A6000 GPU. We employ an SGD optimiser for 60 epochs using the Cosine-Warm-Restart [35] as the learning rate scheduler. A batch size of 128 (STL10 is set to 32), an initial learning rate of 0.05, a momentum of 0.5, and a weight decay of 0.001 are used.

Comparison Methods. We compare with three traditional perception-centred methods: MedianCut [23], MedianCut+Dither [17] and OCTree [19], and a task-centred CNN-based method ColorCNN [25]. Specifically, for ColorCNN, we adopt the same training strategy as we did for CQFormer. For the traditional methods, we conduct comparative experiments as described in [25].

4.2. Object Detection Task Evaluation

Table. 1 shows the object detection results on MS COCO dataset [34] using Sparse-RCNN [43] detector. Our CQFormer outperforms all other methods in terms of AP value performance under all colour quantisation levels, from 1-bit to 6-bit. This substantial improvement demonstrates the effectiveness of our CQFormer in the object detection task. We also evaluate CQFormer with other popular detectors in Supplementary Material.

4.3. Classification Task Evaluation

Results. Fig. 4 presents comparisons to the state-of-the-art methods on the four datasets. Our proposed CQFormer (solid blue line) has a consistent and obvious improvement over all other methods in extremely low-bit colour space (less than 3-bit). Moreover, our CQFormer archives a superior performance than the task-centred method ColorCNN [25] under all colour quantisation levels from 1-bit to 6-bit.

Discussion of classification under a large colour space. Similar to the task-centred ColorCNN, our classification performance is also inferior to the traditional method under a large colour space (greater than 4-bit), despite superior performance on object detection task. As extensively discussed in [25], this is a common characteristic for non-clustering based quantisation.

Here we are satisfied with our classification performance under the small colour space for two reasons: First of all,




Figure 4: Top-1 classification accuracy of colour-quantised images on four datasets with Resnet-18 [21] networks.

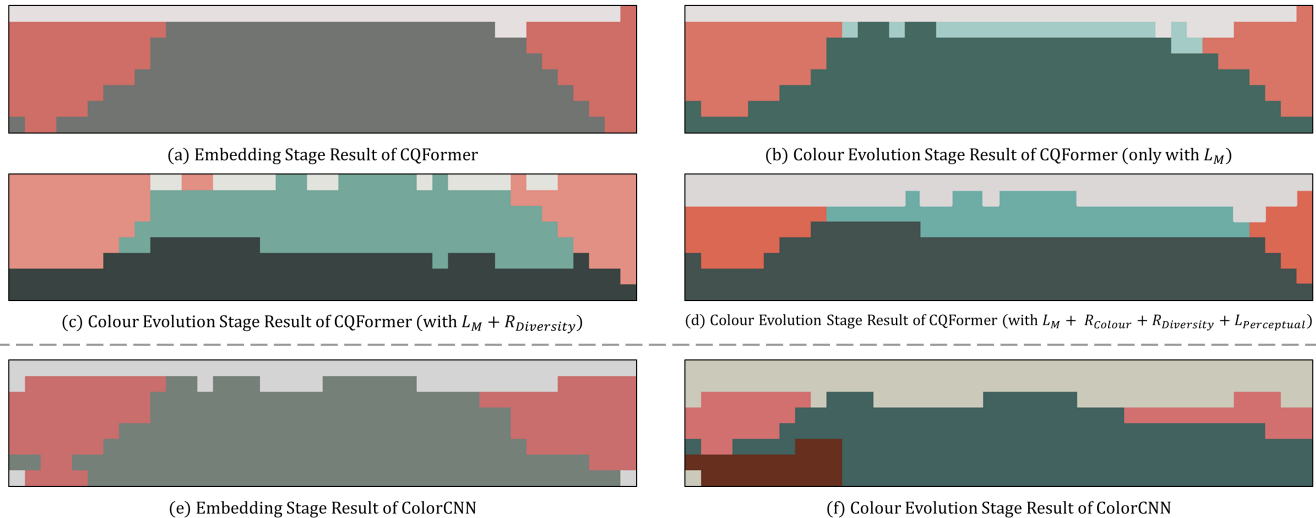


Figure 5: (a) is the embedding stage result of CQFormer. (b) (c) (d) are the colour evolution stage results of CQFormer with different combinations. (e) is the embedding stage result of ColorCNN [25]. (f) is the colour evolution stage result of ColorCNN [25].

most human languages only use fewer than 3-bit colour terms. This implies that discovering more colours not only compromises the principle of efficiency but also goes contrary to the expectation of better perceptual effects. In other words, the performance of the CQFormer on limited colour categories may hint at the optimal outcome restricted by the unique neurological structure of human vision and cognition, which are reflected in a wide array of languages in turn.

Secondly, larger colour space naturally preserves more visual fidelity. 6-bit colour space could already deliver a vivid image, and the recognition performance on these images is only marginally below the original image. It could neither save much storage space nor reveal special knowledge. Therefore, we only focus on the small colour space.

4.4. Colour Evolution Evaluation

Settings. This experiment is based on the classification task on CIFAR10 [28] dataset using the Resnet-18 [21] classifier. In the embedding stage, we embed the 1978 Nafaanra three-colour system (Fig. 3 (c)) using the colour probability map embedding and set C as 3. The reason why we

choose Nafaanra is analysed in the Supplementary Material. Here we set $\tau = 1.0$, and force the fourth colour to be split from light (‘fiNge’), dark (‘wOO’), and warm or red-like (‘nyiE’), respectively. We optimise the loss function $\mathcal{L}_{\text{Full-Embedding}}$ in Eq. 14 for the initial 40 epochs. In the colour evolution stage, we design various combinations of loss function, *i.e.* only \mathcal{L}_M , $\mathcal{L}_M + R_{\text{Diversity}}$ and $\mathcal{L}_M + R_{\text{Colour}} + R_{\text{Diversity}} + \mathcal{L}_{\text{Perceptual}}$, and minimise them separately for the subsequent 20 epochs without the restriction of colour size. Additionally, we also conduct a colour evolution experiment on the *task-centred* ColorCNN [25] for comparison, which is slightly modified to implement the same embedding and evolution details.

Results and Analysis. The results of colour evolution evaluation is shown in Fig. 5. First, during the embedding stage, the WCS colour probability map generated by the CQFormer (Fig. 5 (a)) is similar to Nafaanra (Fig.3 (c)). Therefore, we successfully embed the colour naming system of human language into the latent space of CQFormer. Second, during the colour evolution stage, as shown in Fig. 5 (b)(c)(d), the CQFormer automatically evolves the fourth colour that

is split from dark ('wOO') and close to yellow-green under *all combinations*, matching the basic colour terms theory [27]. Third, although the ColorCNN also evolves the fourth colour close to brown (Fig. 5 (f)), it fails to match the basic colour terms theory [27] since the yellow-green and blue are skipped.

We are not able to see the fourth colour split from either light ('fiNge') or warm/red-like ('nyiE') in the WCS colour probability map, since only 3.7% (if split from light ('fiNge')) or 5.9% (if split from warm/red-like ('nyiE')) of all pixels are assigned to the fourth colour, compared with 23.5% (split from dark ('wOO')). Very interestingly, this phenomenon echoes the evolution of the information bottleneck (IB) colour naming systems [50], where the fourth colour should be split from dark in the "dark-light-red" colour scheme.

Although the CQFormer can also evolve the fourth colour automatically when optimising \mathcal{L}_M alone, it only covers a small portion of the WCS probability map in Fig. 5 (b). As the items of the perceptual structure loss increase, *i.e.* from Fig. 5 (b) to (d), these four colours have clearer borders, a more logical proportion and are each more internally clustered. This suggests that the colour evolution is not fully complete when considering machine accuracy alone. In other words, thanks to the CQFormer's ability to integrate the need for both machine accuracy and human perception, the discovered colour naming system evolves more thoroughly and effectively, which is mirrored in colour naming patterns of human language [49].

5. Limitation and Discussion

While the complexity-accuracy trade-off of machine-discovered colour terms, as shown in Fig. 1 (b), is quite similar to the numerical limit of categorical counterparts for human languages, the current work is still preliminary. As shown in Fig. 1, the newly discovered WCS colour probability map is still quite different from the human one. A more accurate language evolution replication needs to consider more complex variables such as environmental idiosyncrasies, cultural peculiarities, functional necessities, technological maturity, learning experience, and intercultural communication.

Another promising direction would be associating the discovered colours with human colour terms. This would involve much research on Nature Language Processing, and we hope to discuss it with experts from different disciplines in future works. Last but not least, the AI simulation outcome contributes to the long-standing universalist-relativist debate of the linguistic community on colour cognition. Though not entirely excluding the cultural specificity of the colour schemes, the machine finding strongly supports the universalist view that an innate, physiological principle constraints, if not determine, the evolutionary sequence and distributional

possibilities of basic colour terms in communities of different cultural traditions. The complexity-efficiency principle is confirmed by the finding that the numerical limitation of colour categories could lead to superior performance on colour-specific tasks, contrary to the intuitive expectation that complexity breeds perfection. The independent AI discovery of the "green-yellow" category on the basis of the fundamental tripartite "dark-light-red" colour scheme points to the congruence of neural algorithms and human cognition and opens a new frontier to test contested hypothesis in the social sciences through machine simulation. We would be more than delighted if this tentative attempt would prove to be a bridge to link scholars of different disciplines for more collaboration and generate more fruitful results.

Acknowledgement

This work was supported by JST Moonshot R&D Grant Number JPMJMS2011 and JST ACT-X Grant Number JPM-JAX190D, Japan.

References

- [1] Hyojin Bahng, Seungjoo Yoo, Wonwoong Cho, David Keetae Park, Ziming Wu, Xiaojuan Ma, and Jaegul Choo. Coloring with words: Guiding image colorization through text-based palette generation. In *Proceedings of the european conference on computer vision (eccv)*, pages 431–447, 2018. 3
- [2] B. Berlin and P. Kay. *Basic Color Terms: Their Universality and Evolution*. University of California Press, 1969. 1, 2
- [3] Thomas Boutell. Png (portable network graphics) specification version 1.0. Technical report, 1997. 3
- [4] Federico Camposeco, Andrea Cohen, Marc Pollefeys, and Torsten Sattler. Hybrid scene compression for visual localization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7653–7662, 2019. 3
- [5] Nicolas Carion, Francisco Massa, Gabriel Synnaeve, Nicolas Usunier, Alexander Kirillov, and Sergey Zagoruyko. End-to-end object detection with transformers. In *European conference on computer vision*, pages 213–229. Springer, 2020. 5
- [6] Rahma Chaabouni, Eugene Kharitonov, Emmanuel Dupoux, and Marco Baroni. Communicating artificial neural networks develop efficient color-naming systems. *Proceedings of the National Academy of Sciences*, 118(12):e2016569118, 2021. 3
- [7] Huiwen Chang, Ohad Fried, Yiming Liu, Stephen DiVerdi, and Adam Finkelstein. Palette-based photo recoloring. *ACM Trans. Graph.*, 34(4):139–1, 2015. 3
- [8] Junho Cho, Sangdoo Yun, Kyoung Mu Lee, and Jin Young Choi. Palettenet: Image recolorization with given color palette. In *Proceedings of the IEEE conference on computer vision and pattern recognition workshops*, pages 62–70, 2017. 3
- [9] François Chollet. Xception: Deep learning with depthwise separable convolutions. In *Proceedings of the IEEE con-*

- ference on computer vision and pattern recognition*, pages 1251–1258, 2017. 5
- [10] Adam Coates, Andrew Ng, and Honglak Lee. An analysis of single-layer networks in unsupervised feature learning. In *Proceedings of the fourteenth international conference on artificial intelligence and statistics*, pages 215–223. JMLR Workshop and Conference Proceedings, 2011. 7
- [11] Ziteng Cui, Kunchang Li, Lin Gu, Shenghan Su, Peng Gao, ZhengKai Jiang, Yu Qiao, and Tatsuya Harada. You only need 90k parameters to adapt light: a light weight transformer for image enhancement and exposure correction. In *33rd British Machine Vision Conference 2022, BMVC 2022, London, UK, November 21-24, 2022*. BMVA Press, 2022. 5
- [12] Jelmer P de Vries, Arash Akbarinia, Alban Flachot, and Karl R Gegenfurtner. Emergent color categorization in a neural network trained for object recognition. *Elife*, 11:e76472, 2022. 3
- [13] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei. ImageNet: A Large-Scale Hierarchical Image Database. In *CVPR09*, 2009. 7
- [14] Yining Deng, Charles Kenney, Michael S Moore, and BS Manjunath. Peer group filtering and perceptual color image quantization. In *1999 IEEE International Symposium on Circuits and Systems (ISCAS)*, volume 4, pages 21–24. IEEE, 1999. 3
- [15] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit, and Neil Houlsby. An image is worth 16x16 words: Transformers for image recognition at scale. *ICLR*, 2021. 5
- [16] R. W. Floyd and L. Steinberg. An adaptive algorithm for spatial gray scale. *Proceedings of the Society for Information Display*, 17, 1975. 3
- [17] R. W. Floyd and L. Steinberg. An adaptive algorithm for spatial grayscale. *Proceedings of the Society for Information Display*, 17, 1976. 2, 7
- [18] David A Forsyth and Jean Ponce. *Computer vision: a modern approach*. prentice hall professional technical reference, 2002. 2
- [19] Michael Gervautz and Werner Purgathofer. A simple method for color quantization: Octree quantization. In *New trends in computer graphics*, pages 219–231. Springer, 1988. 2, 3, 7
- [20] Edward Gibson, Richard Futrell, Julian Jara-Ettinger, Kyle Mahowald, Leon Bergen, Sivalogeswaran Ratnasingam, Mitchell Gibson, Steven T Piantadosi, and Bevil R Conway. Color naming across languages reflects color use. *Proceedings of the National Academy of Sciences*, 114(40):10785–10790, 2017. 3
- [21] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016. 5, 7, 8
- [22] Yu-Lin He, Xiao-Liang Zhang, Wei Ao, and Joshua Zhexue Huang. Determining the optimal temperature parameter for softmax function in reinforcement learning. *Applied Soft Computing*, 70:80–85, 2018. 4, 5
- [23] Paul Heckbert. Color image quantization for frame buffer display. *ACM Siggraph Computer Graphics*, 16(3):297–307, 1982. 2, 3, 7
- [24] Dan Hendrycks and Kevin Gimpel. Gaussian error linear units (gelus). *arXiv preprint arXiv:1606.08415*, 2016. 5
- [25] Yunzhong Hou, Liang Zheng, and Stephen Gould. Learning to structure an image with few colors. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10116–10125, 2020. 2, 3, 6, 7, 8
- [26] Paul Kay, Brent Berlin, Luisa Maffi, William R Merrifield, and Richard Cook. *The world color survey*. CSLI Publications Stanford, CA, 2009. 1, 3
- [27] Paul Kay and Chad K McDaniel. The linguistic significance of the meanings of basic color terms. *Language*, 54(3):610–646, 1978. 9
- [28] Alex Krizhevsky, Geoffrey Hinton, et al. Learning multiple layers of features from tiny images. 2009. 2, 7, 8
- [29] Ya Le and Xuan Yang. Tiny imagenet visual recognition challenge. *CS 231N*, 7(7):3, 2015. 1, 7
- [30] Boyi Li, Serge Belongie, Ser-nam Lim, and Abe Davis. Neural image recolorization for creative domains. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2226–2230, 2022. 3
- [31] Ke Li, Shijie Wang, Xiang Zhang, Yifan Xu, Weijian Xu, and Zhuowen Tu. Pose recognition with cascade transformers. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1944–1953, 2021. 5
- [32] DianLong Liang, Lizhi Cheng, and Zenghui Zhang. General construction of wavelet filters via a lifting scheme and its application in image coding. *Optical Engineering*, 42(7):1949–1955, 2003. 3
- [33] Tsung-Yi Lin, Piotr Dollár, Ross Girshick, Kaiming He, Bharath Hariharan, and Serge Belongie. Feature pyramid networks for object detection. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2117–2125, 2017. 7
- [34] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C. Lawrence Zitnick. Microsoft coco: Common objects in context. In David Fleet, Tomas Pajdla, Bernt Schiele, and Tinne Tuytelaars, editors, *Computer Vision – ECCV 2014*, pages 740–755, Cham, 2014. Springer International Publishing. 7
- [35] Ilya Loshchilov and Frank Hutter. Sgdr: Stochastic gradient descent with warm restarts. *arXiv preprint arXiv:1608.03983*, 2016. 7
- [36] Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. In *International Conference on Learning Representations*, 2019. 7
- [37] M.T. Orchard and C.A. Bouman. Color quantization of images. *IEEE Transactions on Signal Processing*, 39(12):2677–2690, 1991. 3
- [38] C Alejandro Parraga and Arash Akbarinia. Nice: A computational solution to close the gap from colour perception to colour categorization. *PloS one*, 11(3):e0149538, 2016. 3
- [39] Charles Poynton. *Digital video and HD: Algorithms and Interfaces*. Elsevier, 2012. 3

- [40] Mohammad Rastegari, Vicente Ordonez, Joseph Redmon, and Ali Farhadi. Xnor-net: Imagenet classification using binary convolutional neural networks. In Bastian Leibe, Jiri Matas, Nicu Sebe, and Max Welling, editors, *Computer Vision – ECCV 2016*, pages 525–542, Cham, 2016. Springer International Publishing. 2
- [41] Claude Elwood Shannon. A mathematical theory of communication. *The Bell system technical journal*, 27(3):379–423, 1948. 2
- [42] Katarzyna Siuda-Krzywicka, Christoph Witzel, Emma Chabani, Myriam Taga, Cecile Coste, Noella Cools, Sophie Ferrieux, Laurent Cohen, Tal Seidel Malkinson, and Paolo Bartolomeo. Color categorization independent of color naming. *Cell reports*, 28(10):2471–2479, 2019. 3
- [43] Peize Sun, Rufeng Zhang, Yi Jiang, Tao Kong, Chenfeng Xu, Wei Zhan, Masayoshi Tomizuka, Lei Li, Zehuan Yuan, Changhu Wang, and Ping Luo. Sparse r-cnn: End-to-end object detection with learnable proposals. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 14454–14463, 2021. 7
- [44] Jeya Maria Jose Valanarasu and Vishal M Patel. Unext: Mlp-based rapid medical image segmentation network. *arXiv preprint arXiv:2203.04967*, 2022. 4
- [45] Joost Van De Weijer, Cordelia Schmid, and Jakob Verbeek. Learning color names from real-world images. In *2007 IEEE conference on computer vision and pattern recognition*, pages 1–8. IEEE, 2007. 3
- [46] Xiaolin Wu. Color quantization by dynamic programming and principal analysis. *ACM Transactions on Graphics (TOG)*, 11(4):348–372, 1992. 3
- [47] Jiwei Yang, Xu Shen, Jun Xing, Xinmei Tian, Houqiang Li, Bing Deng, Jianqiang Huang, and Xian-sheng Hua. Quantization networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7308–7316, 2019. 2
- [48] Seungjoo Yoo, Hyojin Bahng, Sunghyo Chung, Junsoo Lee, Jaehyuk Chang, and Jaegul Choo. Coloring with limited data: Few-shot colorization via memory augmented networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11283–11292, 2019. 3
- [49] Noga Zaslavsky, Karee Garvin, Charles Kemp, Naftali Tishby, and Terry Regier. The evolution of color naming reflects pressure for efficiency: Evidence from the recent past. *Journal of Language Evolution*, 2022. 1, 2, 6, 9
- [50] Noga Zaslavsky, Charles Kemp, Terry Regier, and Naftali Tishby. Efficient compression in color naming and its evolution. *Proceedings of the National Academy of Sciences*, 115(31):7937–7942, 2018. 9
- [51] Noga Zaslavsky, Charles Kemp, Naftali Tishby, and Terry Regier. Color naming reflects both perceptual structure and communicative need. *Topics in Cognitive Science*, 11(1):207–219, 2019. 2
- [52] Noga Zaslavsky, Charles Kemp, Naftali Tishby, and Terry Regier. Communicative need in color naming. *Cognitive Neuropsychology*, 2019. 3