



Not just a language with white faces: Analysing #taalmonument on Instagram using machine learning

**Authors:**

Eduan Kotze¹ 
Burgert Senekal¹ 

Affiliations:

¹Department of Computer Science and Informatics, Faculty of Natural and Agricultural Sciences, University of the Free State, Bloemfontein, South Africa

Corresponding author:

Burgert Senekal,
burgertsenekal@yahoo.co.uk

Dates:

Received: 29 Apr. 2020
Accepted: 18 Sept. 2020
Published: 15 Dec. 2020

How to cite this article:

Kotze, E. & Senekal, B., 2020, 'Not just a language with white faces: Analysing #taalmonument on Instagram using machine learning', *The Journal for Transdisciplinary Research in Southern Africa* 16(1), a871. <https://doi.org/10.4102/td.v16i1.871>

Copyright:

© 2020. The Authors.
Licensee: AOSIS. This work is licensed under the Creative Commons Attribution License.

Read online:

Scan this QR code with your smart phone or mobile device to read online.

From the late 19th century, and especially during apartheid (1948–1994), Afrikaans became inextricably tied with white people, white domination and apartheid. This association has persisted after 1994, and calls to preserve Afrikaans are often derided with claims that the preference for Afrikaans is also a preference for racial segregation. In such anti-Afrikaans views, Afrikaans is seen as synonymous with white people and apartheid despite the fact that Afrikaans was never exclusively spoken by white people. This prejudice towards Afrikaans is also shown towards the Afrikaanse Taalmonument, which was unveiled in 1975 to commemorate this language.

Using machine learning and a large dataset of social media posts from Instagram, we show that not only white people visit this monument to Afrikaans, but also take pictures here and post about it on one of the largest social media platforms. As such, we show that the interest in this monument – just like the language itself – is not exclusively tied to one race. We also make suggestions for further research, such as using machine learning for image recognition using social media datasets that could illuminate how other South African monuments are seen in the contemporary world.

Keywords: Afrikaans; Taalmonument; machine learning; convolutional neural networks; Instagram; apartheid.

Introduction

Since the late 19th century, Afrikaans 'was constructed as a "white language," with a "white history" and "white faces"' (Willemse 2017). Because the Afrikaner-dominated National Party (NP) carried out its policy of racial segregation (apartheid) in South Africa from 1948 to 1994, Afrikaans also became associated with apartheid. In particular, the 1976 Soweto riots, which was to a large extent opposition towards Afrikaans as a medium of education, turned the focus of anti-apartheid resistance towards Afrikaans, "This rebellion stigmatized or "further stigmatized" Afrikaans, because the apartheid policy and its application caused injustice and increased a negative attitude towards Afrikaners and Standard Afrikaans' (Steyn 2014:418).¹

Afrikaans is still associated with the apartheid government and related concepts such as oppression and the restriction of freedom, which has led to a resentment towards the language by a large proportion of the South African population (Van Zyl & Rossouw 2016:310). Recent protests at South African university campuses (e.g. #AfrikaansMustFall and #OpenStellenbosch) saw black students mobilising to remove Afrikaans as a language of tertiary education, arguing that it remained a barrier to education, offered an unfair advantage to white students, perpetuated racial segregation and alienated black students. This hostility towards Afrikaans can also be seen in the conduct of African National Congress (ANC) officials, in particular Gauteng MEC for Education, Panyaza Lesufi, and Minister of Higher Education, Blade Nzimande, who have made numerous statements against Afrikaans (Friedman 2019; Steward 2014). Nzimande, for instance, called the private Afrikaans-only tertiary education institution, Akademia, 'racist' because of its language policy (Steward 2014), whilst Lesufi made similar comments about Sol-Tech (Friedman 2019). Such views ignore the fact that the majority of Afrikaans speakers (60%) are not white people (Willemse 2017) but have become commonplace in South Africa nevertheless.

The monument to the Afrikaans language, the Afrikaanse Taalmonument (Afrikaans Language Monument), is likewise associated with apartheid by some, and there have been calls to dismantle

1. Author's own translation from the original Afrikaans, 'Hierdie opstand het Afrikaans gestigmatiseer of "verder gestigmatiseer," want die apartheidsbeleid en die toepassing daarvan het onreg veroorsaak en 'n negatiewe gesindheid teenoor Afrikaners en Standaardafrikaans laat toeneem'.

the Taalmonument in the interest of nation-building (Smith 2013:124; Van Zyl & Rossouw 2016:310). Groenewald (2018:230) calls the ANC-regime 'antagonistic to the language that the monument valorises', and hence hostility towards the monument itself can be expected.

The current study investigates posts made with the hashtag, #taalmonument, on the social media platform, Instagram. As Instagram posts constitute a voluntary association with this monument in the public sphere, the objective of the current study is to determine whether only white people – the race associated with Afrikaans – voluntarily associate themselves with this monument or whether people of other races do the same and to what extent. To this end, we develop, train and evaluate our own machine learning image recognition classifier after constructing our own annotated corpus of images, which is also benchmarked against an internationally recognised dataset. We also make suggestions for future research.

Background to the Taalmonument

The first proposal to erect a monument to Afrikaans was made at a commemoration of the founding of the *Genootskap van Regte Afrikaners* (Association of Real Afrikaners) in 1942 (Groenewald 2018:227; Van Zyl & Rossouw 2016:299). Following this proposal, the Afrikaanse Taalmonumentkomitee (Afrikaans Language Monument Committee) was founded to raise funds for this purpose (Groenewald 2018:227; Van Zyl & Rossouw 2016:299). More than 20 years later, in 1964, a competition was held to select an architect to design the monument, and the architect Jan Van Wijk was chosen to design the monument (De Vaal-Senekal, De Kock & Putter 2018:198; Van Zyl & Rossouw 2016:300). The monument was unveiled by Prime Minister BJ Vorster on 10th October 1975, and the accompanying Taalmuseum (Language Museum) was inaugurated on 14th August 1975 (De Vaal-Senekal et al. 2018:197; Van Zyl & Rossouw 2016:298).

A monument to Afrikaans will inevitably be placed in the racialised discourse that is associated with this language. Although Afrikaans is currently associated with white people and apartheid, this was not always the case: When the *Genootskap van Regte Afrikaners* was founded in 1875, most Afrikaans speakers were not white people, and Afrikaans was often referred to as a *hotnotstaal*² (Groenewald 2018:228). As Willemsse reminds us, 'Afrikaans also has a "black history" rather than just the known hegemonic apartheid history inculcated by white Afrikaner Christian national education, propaganda and the media'. Throughout the apartheid years (1948–1994), however, the Afrikaner was depicted as a white nation, with Afrikaans-speaking coloured people marginalised by the apartheid state. Today, the majority of Afrikaans speakers are not white people, and while 60% of South African white people have Afrikaans as a first language, over 90% of the coloured population

²A term for a person of colour from the Cape area in South Africa. The use of this term is now deprecated and considered offensive. The term 'hotnot' was historically used to refer to the non-Bantu indigenous nomadic pastoralist people of the Western Cape Province of South Africa. The preferred name for the non-Bantu indigenous people is currently Khoi, Khoikhoi or Khoisan.

speak Afrikaans as a first language (Smith 2013:133). Nevertheless, the Afrikaner is generally seen as a white nation (Senekal 2019) (note that *Afrikaner* and *Afrikaans-speaking* are two different labels, the former generally denoting an ethnic group and the latter a linguistic group). In light of this association between white people and Afrikaans, calls for the preservation of Afrikaans are often seen as an attempt to maintain segregation and 'white privilege' (see e.g. Pilane 2015).

The Taalmonument symbolises Afrikaans's diverse roots, including Western European (Dutch, French, German and Portuguese), Malaysian and African languages, including those of the Khoi-Khoi, San and other black Africans (Smith 2013:144; Van Wijk 2014:76; Van Zyl & Rossouw 2016:300–301). Nevertheless, there has been fierce criticism of this monument, including that it is an 'apartheidmonument' (Van Zyl & Rossouw 2016:309). However, Van Wijk (2014:21) states that he did not design the monument for white Afrikaners but rather for the language itself (see also Van Zyl & Rossouw 2016:310). Moreover, an effort was made to secure the attendance of coloured Afrikaans speakers and authors at the inauguration of the Taalmonument in 1975 (Smith 2013:146; Van Zyl & Rossouw 2016:309). A poem by Adam Small (one of the most prominent coloured Afrikaans authors), 'Nkosi sikelel' iAfrika', was also recited at the opening (Smith 2013:146). From the beginning, then, the Taalmonument has aimed at shedding the stigma of Afrikaans being a language reserved for white people. However, with the Soweto riots occurring just the year after the opening of the Taalmonument, this attempt at making Afrikaans more inclusive seems to have had little effect.

Today, the Taalmonument still aims at inclusivity, 'The ATM strives for all South Africans to appreciate Afrikaans. In this spirit, the ATM works hard to encourage and support Afrikaans among the youth and non-mother-tongue speakers' (De Vaal-Senekal et al. 2018:198, see also Van Zyl & Rossouw 2016:311; Smith 2013:138).

This effort to broaden the appeal of the Taalmonument and the museum should lead to a diverse collection of visitors. In the contemporary world, visitors to monuments and museums often share their visits with others on social media platforms, such as Instagram, which provides the opportunity to analyse social media posts to obtain a better understanding of who visits monuments and why. The following section provides a short background on Instagram.

Instagram

Being founded in 2010, Instagram quickly became a major role player as a social media platform. Currently, Instagram has around a billion worldwide users each month and 500 million users each day, with over 50 billion photos uploaded to date (Aslam 2020). In South Africa, Facebook is the most popular social media platform, followed by YouTube, WhatsApp, Facebook Messenger, LinkedIn, Twitter

and Instagram (Qwerty 2017:12). Instagram is a photo-based platform that allows only photo and video posts, that is, no text-only posts similar to Facebook and Twitter.

Instagram is, however, not representative of the entire population of a country as Instagram users tend to be younger (Anderson & Jiang 2018; Aslam 2020; Duncan 2016). This is particularly relevant in the current study, as people who visit the Taalmonument and post pictures of their visits later will probably be from a younger generation that is less tied to a first-hand experience of apartheid and the NP. Note, however, that we do not have access to users' ages.

To investigate whether only white people or people of different races associate themselves with the Taalmonument on Instagram, we first had to train a model to distinguish between different races. The following section provides a background to machine learning for racial classification, after which we discuss the specific methods we used.

Machine learning for image classification

Machine learning is a subfield of artificial intelligence (AI) and was developed from the 1960s onwards (Kononenko 2001; Michie 1968), in particular through the works of Rosenblatt (1962), Nilsson (1965) and Hunt, Martin and Stone (1966). The field gained ground in the most recent two decades because of the big data revolution (Jordan & Mitchell 2015:256), leading Jordan and Mitchell (2015:260) to claim, 'machine learning is likely to be one of the most transformative technologies of the 21st century'.

A large amount of recent research has been directed towards identifying race in images using machine learning (Fu, He & Hou 2014; Trivedi & Amali 2017; Vo, Nguyen & Le 2018). Although the concept of race is a contentious issue, particularly as the term is often used interchangeably or confused with ethnicity (see, e.g. Bartlett 2001; Collins 2004; Markus 2008), Fu et al. (2014:2483) define the difference between race and ethnicity simply, 'race refers to a person's physical appearance or characteristics, while ethnicity is more viewed as a culture concept, relating to nationality, rituals and cultural heritages, or even ideology'. We prefer this simple distinction between race and ethnicity and focus the rest of our discussion on race.

Racial classification is in one sense a highly controversial topic, because it carries the baggage of the Population Registration Act (Union of South Africa 1950) that, along with other apartheid-era legislation, led to racial discrimination and human rights abuses in South Africa before 1994. In contrast, racial classification is not controversial in contemporary South Africa: Broad-Based Black Economic Empowerment (BBBEE), as well as the discourse around *white monopoly capital*, *transformation*, *white privilege* and *land expropriation*, assumes racial categories. Despite the abolishment of racial categories in South Africa during the final years of apartheid, racial categories have persisted in

the South African census and in public discourses. Most university staff have experienced being obligated to indicate their race on administrative forms as well, with racial categories reminiscent of the Population Registration Act (Union of South Africa 1950) (white-, black-, coloured-, Indian people and other). We would therefore like to emphasise that we trained a model to conduct racial classification because the discourse on Afrikaans and the Taalmonument is already racialised; the irony of deracialising this discourse is that we first need to be able to distinguish between races to ascertain whether visitors to the Taalmonument who post about their visits afterwards on Instagram belong to one or various races.

A variety of racial classification methods using machine learning have been proposed. Fu et al. (2014:2487) note 'statistically significant variances in facial anthropometric dimensions between all race groups', which 'pave the way of anthropometry-based automatic race recognition'. The question is what to measure. There is a common misconception that race is defined by a skin colour (as exemplified by referring to people as 'white' or 'black'), and numerous efforts have been made to use skin colour to differentiate between races, but Fu et al. (2014:2485) argue 'skin color is such a variable visual feature within any given race that it is actually one of the least important factors in distinguishing between races'. A second view holds that 'physical characteristics such as hairshaft morphologic characteristics and craniofacial measurements are viewed as significant indicators of race belongings' (2014:2485), whilst another method compares the eyes of subjects; Fu et al. (2014:2490) note 'Statistically significant race differences in retinal geometric characteristics', which have been reported in several studies. We opted for a more holistic approach by extracting whole faces and teaching a model to which race faces belong, as discussed below.

Depending on the criteria and level of analysis, there are between three and 200 races (Coon 1962). Fu et al. (2014:2485) distinguish between seven races, which cover about 95% of the world population: African/African American, caucasian, East Asian, Native American/American Indian, Pacific Islander, Asian Indian and Hispanic/Latino. These seven races, of course, exclude coloured people. In adapting racial classifications for the South African context, we initially used the classifications suggested by Jan Raats, whose classification was used by the NP government through the Population Registration Act (James 2012; Union of South Africa 1950) and can still be found on administrative forms in South Africa today. These categories distinguish between four races: white-, black-, coloured- and Asiatic people (we substitute his classification of 'bantu' for the more politically acceptable term 'black'). However, the difference between Indian- and Asian people is so striking that we decided to split the Asiatic category into Asian- and Indian people.

People of a mixed-race origin pose a significant challenge to existing facial recognition models (Fu et al. 2014:2502).

This predicts that there will be difficulty in classifying South Africans who have been mixing for the past 350 years, especially for the coloured population. Afrikaners, although generally considered white people, are also not exclusively caucasian in their genetic makeup (Erasmus, Klingenberg & Greeff 2015; Greeff 2007; H. Heese 1979, 1984; J. Heese 1971).

Our experiments confirmed Fu et al.'s (2014:2502) assertion and encountered substantial difficulty in distinguishing between white-, coloured- and black faces. When all five categories were included, we failed to move beyond an accuracy level of 70%, regardless of how we refined our model. We, therefore, simplified our racial categories to a binary classification, white or black, as the objective of the current study is in any case to determine whether only white people associate themselves with the Taalmonument and whether other races do the same, regardless of which race those people belong to.

The following section describes how the model was constructed and trained.

Methods

Model training

#modelsofinstagram dataset

A random sample of images was downloaded from Instagram to collect sufficient training data that could be used in the construction of a classifier. Images placed on Instagram are already annotated to some degree by placing them with a hashtag, but the hashtag indicates to which discourse the image belongs and not necessarily what the content of the image is. A picture with the hashtag #europeans could, for instance, show the picture of an African slave, as Europeans are known for slavery, but the hashtag does not indicate that the content of the image is a black African. We experimented with various possible hashtags that could be used to construct a labelled dataset, but possible hashtags differed considerably across races: Whilst #blackmodels and #indianmodels collected images of people belonging to these races, #whitemodels had a very limited selection of images and #colouredmodels created problems with the different meanings associated with the term. Hashtags such as #san, #european and #sotho did not return a meaningful number of relevant images. The hashtag #afrikaner delivered a considerable number of irrelevant images, again partly because the term carries different meanings in different languages. We eventually decided to use a single hashtag, #modelsofinstagram, and used an annotator to classify people according to race. The annotator is in his late thirties and thoroughly familiar with racial categories in a South African context.

After the annotator had labelled the images, we began work on developing an image classifier. To classify an image according to race, we first needed to perform face detection and extract a face from an Instagram image, because this reduced the amount of noise in an image. For automatic face detection, we used opencv-python 4.2.0.34 (Heinisuo 2020),

a wrapper package for OpenCV Python bindings to perform image processing. OpenCV is a modern implementation of the novel Classifier Cascade face detection algorithm (Viola & Jones 2001) and provides the *CascadeClassifier* class that allowed us to create a cascade classifier for face detection. A cascade, in machine learning terms, is an approach where a function is trained from numerous positive and negative images. This will allow the image classifier to detect objects (such as faces) in images. The result is that OpenCV allows us to extract faces from images, regardless of how many faces there are in a single image.

Using the face detection classifier, we were able to successfully detect and extract 3534 faces (2129 that were annotated as white people and 1405 that were annotated as black people) from our training dataset, using a confidence factor of 0.98 (in other words, we only allowed OpenCV to extract faces if it was 0.98% certain that it was a face it had identified). We then randomly selected from each image dataset to create the training and testing datasets. Table 1 shows the number of images we used for the training and validation of the model.

UTKFACE dataset

We also wanted to compare our annotator's classification by benchmarking our classifier with an internationally recognised dataset. The UTKFace dataset by Zhang, Song and Qi (2017) is a large face dataset with a long age span (subjects of between 0 and 116 years old) and consists of over 20 000 images with annotations in terms of age, gender and race. We applied an age filter (18–65 years old) on the dataset resulting in 17 655 images, as the #modelsofinstagram facial images were extracted from Instagram users who will most likely fall within this age range, as will the images posted with #taalmonument. We then randomly selected from these images selecting only white people (race = 0) and black people (race = 1) images to create the training and testing datasets. We did not filter on gender, because we wanted our classifier to function across genders. Table 2 shows the training and validation sets we used from the UTKFace dataset.

Data augmentation

As both datasets consist of a relatively small number of training examples, one can inadvertently introduce overfitting into a model. Overfitting occurs when a model learns the noise instead of the signal of the training data and consequently will not generalise well from training data

TABLE 1: Number of training and validation images for dataset 1.

Variables	White people	Black people	Total
Training	1000	1000	2000
Validation or testing	400	400	800
Total	1400	1400	2800

TABLE 2: Number of training and validation images for dataset 2.

Variable	White people	Black people	Total
Training	3000	3000	6000
Validation or testing	500	500	1000
Total	3500	3500	7000

to unseen data. In predictive modelling, the signal is the underlying pattern that the machine learning model should learn from the data. In other words, overfitting refers to when the model does not accurately learn what it is supposed to evaluate because of a small dataset. Suppose a large number of images of dogs also contain cars, the model can mistakenly identify cars with dogs and classify a cat as a dog because of the presence of a car in the image.

One way to overcome overfitting is to introduce *data augmentation* by generating more training examples from the existing training dataset. Data augmentation may include flipping, rotating or blurring images. The goal is to use random transformations that create believable-looking images and consequently artificially increase the number of training examples. For our experiment, we made use of the *ImageDataGenerator* class that is bundled with Keras (Chollet 2017), a Python deep learning library, to create batches of images with real-time data augmentation to both training datasets (`#modelsofinstagram` and `UTKFace`). These included flipping images horizontally, rotating images by 45 degrees and zooming images up to 50% randomly. Finally, we also applied width shift and height shift by factors of 0.15.

Deep learning models


For our machine learning classifiers, we made use of convolutional neural networks (CNNs) (Goodfellow, Bengio & Courville 2016:326). These classifiers are a specialised kind of neural network that process vector space (grid-like topology) datasets. These datasets can be a one-dimensional grid (1-D) such as time-series data or a two-dimensional grid (2-D) of pixels such as image data. Convolutional Neural Networks have been used successfully in applications such as facial recognition and more recently in natural language processing. Examples of CNN image recognition models are MobileNet by Howard et al. (2017), Levi and Hassner's (2015) age and gender recognition model and Campos, Jou and Giró-i-Nieto's (2017) image sentiment recognition model.

Convolutional neural networks consists of a series of *convolutional* and *pooling* layers, and all CNN models have a similar architecture. The architecture of CNN models is shown in Figure 1, which is adapted from Dertat (2017).

As the name convolutional neural network indicates, the neural network model employs a mathematical operation called a convolution. A convolution is a specialised kind of linear operation and enables a CNN to use convolution instead of a general matrix multiplication in at least of one its layers (Goodfellow et al. 2016:327). After a convolution operation, the network will perform pooling to reduce the dimensionality. This enables the network to reduce the number of training parameters and as a result also shortens the training time. The most common type of pooling is max pooling, which is the same type of pooling we use in our classifiers. This enabled the models to reduce the input to the pooling layer (e.g. $32 \times 32 \times 10$ dimensionality) to a $16 \times 16 \times 10$ feature map as illustrated in Figure 2 (adapted from Dertat (2017)).


For our study, we constructed three CNNs: A CNN model consisting of three convolution blocks (Model1), a CNN model consisting of four convolution blocks (Model2) and a CNN model based on the VGG16 model proposed by Simonyan and Zisserman (2014) (Model3).

The first CNN model consists of three convolution blocks (3×3 filter) with the same padding and a max pool layer (2×2 filter) in each of them resulting in eight layers. For the



Source: Adapted from Dertat, A., 2017, *Applied deep learning – Part 4: Convolutional neural networks*, viewed 22 April 2020, from <https://towardsdatascience.com/applied-deep-learning-part-4-convolutional-neural-networks-584bc134c1e2>

FIGURE 2: Convolutional Neural Network pooling.



Source: Adapted from Dertat, A., 2017, *Applied deep learning – Part 4: Convolutional neural networks*, viewed 22 April 2020, from <https://towardsdatascience.com/applied-deep-learning-part-4-convolutional-neural-networks-584bc134c1e2>

Conv, convolution; Pool, pooling operations; FC, fully connected.

FIGURE 1: Convolutional neural network architecture.

classification block, there were two fully connected layers with 512 units on top of the convolution blocks that were activated by a *relu* activation function. In deep learning neural networks, the activation function is responsible for transforming the summed weight input from a node into the activation of the node or output for that node (Brownlee 2019). Popular activation functions include *sigmoid* (or logistic), *tanh* (hyperbolic tangent) or *relu* (rectified linear units). We opted for *relu* as it allows for backpropagation of errors to train our deep learning models (Goodfellow et al. 2016:226). In total, there were 10 904 097 trainable parameters.

The second CNN model consists of four convolution blocks (3×3 filter) with the same padding and a max pool layer (2×2 filter) in each of them resulting in 12 layers. For the classification block, there was a single fully connected layer with 512 units on top of the convolution blocks that were activated by a *relu* activation function. In total, there were 7 595 809 trainable parameters.

The third CNN model was a scaled-down version of the original VGG16 model proposed by Simonyan and Zisserman (2014). The original VGG16 model consists of five convolution blocks (3×3 filter) with a max pool layer (2×2 filter) in each of them where the '16' refers to 16 layers that have weights. Our model consists of four convolution blocks (3×3 filter) with the same padding and a max pool layer (2×2 filter) in each of them resulting in 14 layers that have weights. For the classification block, there were two fully connected layers with 512 units on top of the convolution blocks that were activated by a *relu* activation function. In total, there were 12 790 433 trainable parameters.

All three models output class probabilities based on a binary classification by the *sigmoid* activation function for output. We made use of the *ADAM* optimiser and a *binary cross entropy loss* function. *Adam* is an adaptive learning rate optimisation algorithm specifically designed for deep learning (Kingma & Ba 2017). As we are using a binary classifier, our loss function will also be binary and use cross entropy to measure how far from the true value (0 or 1) our prediction of each image was. The loss function will then average these class-wise errors to obtain the final loss (Peltarion 2020). We also experimented with dropout, a regularization technique used to reduce the overfitting of a network. Dropout takes a fractional number as its input value, in the form such as 0.1, 0.2, 0.4, etc. This means dropping out 10%, 20% or 40% of the output units randomly from the applied layer. For CNN model 1 and CNN model 3, we applied dropout to the last max pool layer (0.3 for Model 1 and 0.2 for Model 3).

Testing the models

We trained the three CNN models on both datasets. As the classifier is a binary classifier (only two labels, i.e. white people or black people), we report the precision, recall, F1 and accuracy as the evaluation metrics used to assess the performance of the CNN models. Precision is the ability of a

classifier not to label a sample as positive if it was negative. Recall is the ability of the classifier to find all the positive samples. Accuracy returns the number of correctly classified samples whilst F1 is the weighted average of precision and recall. As the training of the models took a substantial amount of time, we did not train using *n*-fold cross validation. Cross-validation is a resampling technique to evaluate machine learning models on a limited dataset whilst *n* (in *n*-fold) refers to the number of groups that a given dataset is split into. Instead, we made use of Model Checkpoint and Early Stopping. Model Checkpoint monitors a specific parameter of the model (we used *val_loss* or validation loss) and Early Stopping will stop the training process of the model if there was no improvement in validation loss after a number of epochs. An epoch refers to the number of times that a learning algorithm with work through a training dataset. We set the maximum number of epochs at 100 and allowed the model to stop after 10 epochs if there was no improvement in validation loss. After the training was completed, we tested each model with both #modelsofinstagram (*n* = 800) and UTFFace (*n* = 1000) testing datasets. Table 3 provides the test evaluation metrics of each model and dataset.

From the testing of the models, CNN Model2 with our own #models of instagram dataset performed the best. When examining the model during testing, we noted a test loss of 0.106758 with an accuracy of 0.97. The model reached the optimal training validation loss value at *n* = 33 epochs. In other words, our model is capable of predicting a person's race with 97% accuracy. With the model created, trained and evaluated, we could now apply it to a dataset of images downloaded with the hashtag #taalmonument, as discussed in the following section.

Data gathering

Before we could investigate the race of people that posted with the hashtag #taalmonument, we first had to download all posts tagged with this hashtag. Posts were downloaded using the application, InstaBro, on 14 February 2020. The first post was made on 01 July 2012, meaning that the dataset spans over 7 years. There were 2988 photos posted with this hashtag (#taalmonument) during this period. Note that we could not gather any data about users, including their names, age, location or gender. Importantly, we could only download posts from public profiles, that is, we were not required to follow users in order to include their content in the analysis below. In other words, these posts were made openly, in front of an audience numbering around a billion, which means that the dataset constitutes posts made by people who

TABLE 3: The test evaluation metrics of each model and dataset.

Dataset	CNN	Accuracy	Precision	Recall	F1
#modelsofinstagram	Model1	0.95	0.95	0.95	0.95
	Model2	0.97	0.97	0.97	0.97
	Model3	0.97	0.97	0.97	0.97
UTFFace	Model1	0.94	0.94	0.94	0.94
	Model2	0.95	0.95	0.95	0.95
	Model3	0.94	0.94	0.94	0.94

CNN, Convolutional Neural Network; F1, weighted average of precision and recall.

TABLE 4: Results.

Variable	Black people (label = 0)	White people (label = 1)	Total
CNN Model 2	139	529	688

CNN, Convolutional Neural Network.

openly chose to associate themselves with the Taalmonument. Furthermore, not including any information about users prevents violating user privacy. For the same reason, we cannot provide examples of the racial classifications of users and rather use the results in aggregate.

Results

To perform predictions on the unlabeled facial image dataset (#taalmonument), we deployed the best-performing classification model (CNN Model 2). First, the unlabelled dataset from Instagram (#taalmonument) was preprocessed, which included scaling and extracting human faces. From the 2988 unlabelled photos, 668 human faces were identified and extracted using OpenCV (the rest of the photos were of the monument or the landscape around the Taalmonument). We then passed these facial images to our model as input and received a label as output. Table 4 summarizes the results.

The following section discusses these results.

Discussion

Census data show that most Afrikaans speakers are not white people, but as noted in the section discussing the background, the Taalmonument and Afrikaans are both accused of being exclusively white phenomena. The results in the previous section, however, show that this is not entirely the case in our study. Of 688 faces identified from Instagram posts made with the hashtag #taalmonument, 529 (76.89%) were white people and 139 (20.2%) were black people. As we showed our classifier to predict people's race with 97% accuracy, this shows that 20% of people who chose to associate themselves with the Taalmonument are not white people. The key issue here is voluntary association: whilst people may attend an Afrikaans university based on the geographical location, the availability of transport, limited course options or for other reasons, people who take a photo at a monument and post it to Instagram do so willingly and intentionally. Moreover, taking the time to travel to the monument, taking a picture and posting it on Instagram constitute a significant effort on the part of the user. The fact that a substantial number of people who associate themselves with the Taalmonument are not white shows that this monument does not only garner attention from the white population but rather functions in an inclusive capacity, as intended by Van Wijk.

However, it is unclear why only 20% of the faces we identified are not white people, whilst white people are a minority both in the national population of South Africa and amongst Afrikaans speakers. This over-representation of white people may reflect Instagram user demographics (no data are available on the distribution of Instagram use by race in SA), cultural differences or it may indicate a smaller proportional

interest in the monument. It may, for instance, be that a smaller proportion of coloured people show an interest in this monument than is the case for white people, but we have no data to explain this skewed distribution and other factors could be at play.

Of course, although the above shows a diverse association of people with the Taalmonument on Instagram, this study did not conduct a representative investigation into attitudes towards the Taalmonument. Such a study of attitudes can better be conducted using a large sample of questionnaires or interviews. However, the above does show that, contrary to claims that it is an 'apartheidsmonument', users on Instagram take the time and effort to publicly associate themselves with this monument even if they are not white people.

Conclusion

This article showed that people who visit the Taalmonument and post about their visits later on Instagram are from various racial backgrounds. Contrary to the racialised discourse on Afrikaans in South Africa, our study shows that not only white people take the time and effort to travel to this monument, take pictures and post about it afterwards on Instagram – in other words, voluntarily associate with this monument on a highly public global platform. Our study therefore does not suggest that the Afrikaanse Taalmonument is considered to be a 'white people only' or 'apartheid' monument but rather a monument that has enough significance for people of other races to also take the time and effort to take photos here and post about it on social media.

We only investigated one monument and one factor, namely race. Future studies could apply a similar method to investigate the demographics of visitors who post on social media with relation to other museums and monuments in South Africa, including, for example, considering visitors' age and gender. Social media provides a wealth of data with which to investigate how museums and monuments function in the contemporary world and in people's lives, and much of this opportunity has not been realised in academic research yet.

Acknowledgements

Authors' contributions

All authors contributed equally to this work.

Funding information

University of the Free State Interdisciplinary Research Grant.

Ethical consideration

This article followed all ethical standards for carrying out research.

Data availability statement

The data are not publicly available due to privacy restrictions.

Disclaimer

The views and opinions expressed in this article are those of the authors and do not necessarily reflect the official policy or position of any affiliated agency of the authors.

References

- Anderson, M. & Jiang, J., 2018, *Teens, social media and technology 2018*, viewed 23 March 2020, from <https://www.pewresearch.org/internet/2018/05/31/teens-social-media-technology-2018/>.
- Aslam, S., 2020, *Instagram by the numbers: Stats, demographics and fun facts*, viewed 23 March 2020, from <https://www.omnicoreagency.com/instagram-statistics/>.
- Bartlett, R., 2001, 'Medieval and modern concepts of race and ethnicity', *Journal of Medieval and Early Modern Studies* 31(1), 39–56. <https://doi.org/10.1215/10829636-31-1-39>
- Brownlee, J., 2019, *A gentle introduction to the Rectified Linear Unit (ReLU)*, viewed 22 April 2020, from <https://machinelearningmastery.com/rectified-linear-activation-function-for-deep-learning-neural-networks/>.
- Campos, V., Jou, B. & Giró-i-Nieto, X., 2017, 'From pixels to sentiment: Fine-tuning CNNs for visual sentiment prediction', *Image and Vision Computing* 65, 15–22. <https://doi.org/10.1016/j.imavis.2017.01.011>
- Chollet, F., 2017, *Deep learning with python*, Manning Publications, Shelter Island, NY.
- Collins, F.S., 2004, 'What we do and don't know about "race," "ethnicity," genetics and health at the dawn of the genome era', *Nature* 36(11), 13–15. <https://doi.org/10.1038/ng1436>
- Coon, C., 1962, *The origins of races*, Knopf, New York, NY.
- De Vaal-Senekal, P., De Kock, C. & Putter, M., 2018, 'Challenges in the archives of the Afrikaans Language Museum, Paarl, Western Cape, South Africa: A case study', *Atlanti* 28(1), 195–205. [https://doi.org/10.33700/2670-451X.28.1.195-205\(2018\)](https://doi.org/10.33700/2670-451X.28.1.195-205(2018))
- Dertat, A., 2017, *Applied deep learning – Part 4: Convolutional neural networks*, viewed 22 April 2020, from <https://towardsdatascience.com/applied-deep-learning-part-4-convolutional-neural-networks-584bc134c1e2>.
- Duncan, F., 2016, *So long social media: The kids are opting out of the online public sphere*, viewed 17 September 2018, from <http://theconversation.com/so-long-social-media-the-kids-are-opting-out-of-the-online-public-square-53274>.
- Erasmus, J.C., Klingenberg, A. & Greeff, J.M., 2015, 'Allele frequencies of AVPR1A and MAOA in the Afrikaner population', *South African Journal of Science* 111(7), 1–6. <https://doi.org/10.17159/sajs.2015/20150074>
- Friedman, B., 2019, *Panyaza Lesufi takes a swipe at private Afrikaans university under construction*, viewed 20 April 2020, from <http://www.702.co.za/articles/360944/panyaza-lesufi-takes-a-swipe-at-private-afrikaans-university-under-construction>.
- Fu, S., He, H. & Hou, Z.-G., 2014, 'Learning race from face: A survey', *IEEE Transactions on Pattern Analysis and Machine Intelligence* 36(12), 2483–2509. <https://doi.org/10.1109/TPAMI.2014.2321570>
- Goodfellow, I., Bengio, Y. & Courville, A., 2016, *Deep learning: Machine learning book*, MIT Press Ltd., Cambridge.
- Greeff, J.M., 2007, 'Deconstructing Jaco: Genetic Heritage of an Afrikaner', *Annals of Human Genetics* 71(5), 674–688. <https://doi.org/10.1111/j.1469-1809.2007.00363.x>
- Groenewald, M., 2018, 'An interrogation of the visual rhetoric of South African graphic designer Ernst de Jong (1934–2016)', Unpublished PhD-thesis, University of Pretoria.
- Heese, H., 1979, 'Identiteitsprobleme gedurende die 17de eeu', *Kronos* 1, 27–33.
- Heese, H., 1984, *Groep sonder grense. Die rol en status van die gemengde bevolking aan die Kaap, 1652–1795*, Universiteit van Wes-Kaapland, Bellville.
- Heese, J., 1971, *Die herkoms van die Afrikaner, 1657–1867*, A.A. Balkema, Kaapstad.
- Heinisuo, O.-P., 2020, *Opencv-python 4.2.0.34*, viewed 22 April 2020, from <https://pypi.org/project/opencv-python/>.
- Howard, A.G., Zhu, M., Chen, B., Kalenichenko, D., Wang, W., Weyand, T., Andreetto, M. & Adam, H., 2017, 'MobileNets: Efficient convolutional neural networks for mobile vision applications', viewed n.d., from <https://arxiv.org/pdf/1704.04861.pdf>.
- Hunt, E., Martin, J. & Stone, P., 1966, *Experiments in induction*, Academic Press, New York, NY.
- James, W., 2012, *The strange career of race classification in South Africa*, viewed 04 April 2019, from <http://politicsweb.co.za/news-and-analysis/the-strange-career-of-race-classification-in-south>.
- Jordan, M.I. & Mitchell, T.M., 2015, 'Machine learning: Trends, perspectives, and prospects', *Science* 349(6245), 255–260. <https://doi.org/10.1126/science.aaa8415>
- Kingma, D.P. & Ba, J., 2017, 'Adam: A method for stochastic optimization', *arXiv Preprint:arXiv:1412.6980v9*.
- Kononenko, I., 2001, 'Machine learning for medical diagnosis: History, state of the art and perspective', *Artificial Intelligence in Medicine* 23(1), 89–109. [https://doi.org/10.1016/S0933-3657\(01\)00077-X](https://doi.org/10.1016/S0933-3657(01)00077-X)
- Levi, G. & Hassner, T., 2015, *Age and gender classification using convolutional neural networks*, s.n., Boston, MA.
- Markus, H.R., 2008, 'Pride, prejudice, and ambivalence: Toward a unified theory of race and ethnicity', *American Psychologist* 63(8), 651–670. <https://doi.org/10.1037/0003-066X.63.8.651>
- Michie, D., 1968, '"Memo" functions and machine learning', *Nature* 218(5136), 19–22. <https://doi.org/10.1038/218019a0>
- Nilsson, N., 1965, *Learning machines*, McGraw-Hill, New York, NY.
- Peltarion, 2020, *Binary crossentropy*, viewed 22 April 2020, from <https://peltarion.com/knowledge-center/documentation/modeling-view/build-an-ai-model/loss-functions/binary-crossentropy>.
- Pilane, P., 2015, *Black students' dissent on afrikaans campuses is about more than just language policy*, viewed 20 April 2020, from <https://www.thedailyvox.co.za/black-students-dissent-on-afrikaans-campuses-is-about-more-than-just-language-policy/>.
- Qwerty, 2017, *The digital landscape in South Africa 2017. A data driven look at South Africa's relationship with digital*, viewed 19 March 2020, from <http://qwertydigital.co.za/wp-content/uploads/2017/08/Digital-Statistics-in-South-Africa-2017-Report.pdf>.
- Rosenblatt, F., 1962, *Principles of neurodynamics*, Spartan Books, Washington, DC.
- Senekal, B.A., 2019, 'Ras en afrikanerretnisiteit: 'n kwantitatiewe ondersoek na huidige opvattinge', *Ensovoort* 40(8), 1.
- Simonyan, K. & Zisserman, A., 2014, *Very deep convolutional networks for large-scale image recognition*, s.n., s.l.
- Smith, S.J., 2013, 'Monumentalising language: Visitor experience and meaning making at the Afrikaanse Taalmonument', Unpublished PhD-thesis, Southern Cross University.
- Steward, D., 2014, *Nzimande and the Afrikaners*, viewed 20 April 2020, from <https://www.politicsweb.co.za/opinion/nzimande-and-the-afrikaners>.
- Steyn, J., 2014, 'Ons gaan 'n taal maak' Afrikaans sedert die Patriot-jare, Kraal Uitgewers, Pretoria.
- Trivedi, A. & Amali, D.G.B., 2017, 'A comparative study of machine learning models for ethnicity classification', *IOP Conference Series: Materials Science and Engineering* 263(4), 042091. <https://doi.org/10.1088/1757-899X/263/4/042091>
- Union of South Africa, 1950, *Population registration act, 1950*, Union of South Africa, Cape Town.
- Van Wijk, J., 2014, *Taalmonument*, Historical Media, Tokai.
- Van Zyl, A. & Rossouw, J., 2016, 'Die Afrikaanse Taalmuseum en -monument in die Paarl: 40 jaar later', *Tydskrif vir Geestewetenskappe* 56(2), 295–313. <https://doi.org/10.17159/2224-7912/2016/v56n2-1a2>
- Viola, P. & Jones, M., 2001, *Rapid object detection using a boosted cascade of simple features*, s.n., I-I, Kauai.
- Vo, T., Nguyen, T. & Le, C.T., 2018, 'Race recognition using deep convolutional neural networks', *Symmetry* 10(564), 1–15. <https://doi.org/10.3390/sym10110564>
- Willemsse, H., 2017, *More than an oppressor's language: Reclaiming the hidden history of Afrikaans*, viewed 20 April 2020, from <https://theconversation.com/more-than-an-oppressors-language-reclaiming-the-hidden-history-of-afrikaans-71838>.
- Zhang, Z., Song, Y. & Qi, H., 2017, *Age progression/regression by conditional adversarial autoencoder*, viewed n.d., from <https://arxiv.org/pdf/1702.08423.pdf>.