

Pawpularity Contest: increasing cuteness of cat & dog photos

Lisa Hoek

Radboud University, Nijmegen, Netherlands

lisa.hoek@ru.nl

ABSTRACT

When people look for a new cat or dog online, the first thing they often see is the pet's profile picture. To increase chances of adoption, online adoption platform 'PetFinder.my' wishes to improve their cuteness ranker tool that scores pet photos based on 'cuteness'. In this research, we investigate what a photographer can do to increase their pet's popularity. We confirm former work that the static features age, breed/colour and size already quite influence how soon a pet is adopted. However, we also found features that are manipulable, i.e., certain flooring or objects that showed strong positive or negative correlations with pet's popularity. Our findings can help pet owners to make attractive photos of their pets and increase their chances of adoption.

1 INTRODUCTION

Cats and dogs are part of many people's lives. However, millions of them live in shelters and sometimes live there for many years. Why is it for some animals so difficult to find a new home, and why are some others immediately adopted? For online adoption platforms, one of the reasons might be the pet's profile picture. Cute-looking pets draw people's attention more, receive more user-clicks and have a bigger chance of being adopted [9]. Not only age, breed/colour and sex might influence the popularity of a pet, there could also be some other factors hidden in the photos that we have not taken into account yet. This research is aimed at trying to find other features in pet's profile pictures that increase (or decrease) pet popularity.

In this paper, we will look at 9912 cats and dogs photos and will retrieve labels/descriptions of what is seen in the photos. We will compare the labels of less popular pet photos and higher popular photos. This cats and dogs dataset is taken from the related Kaggle competition 'PetFinder.my - Pawpularity Contest'¹ in which participants are asked to build models that predict pet photos' popularities. This dataset already includes some additional metadata which was the main reason and starting point of this research. Each pet photo also has a Pawpularity score between 1 and 100. This score indicates how popular the pet was on the website and this will be used to evaluate our findings.

Based on previous literature, we will take a look at pet features that are static such as age, breed/colour and size, but we will also examine whether there are some features influencing popularity that lie in the control of the photographer, such as adding a toy, showing the animal in a cage, etc. So, instead of building our own model that predicts popularity for pet photos (the goal of the competition), we focus our research more towards extracting actual

knowledge by using a multimedia mining approach. We formulate our research question as follows:

Research question. "What features in photos make cats and dogs more 'cute-looking' and thus sooner adopted?"

To answer this research question, we will look at the correlation between the given popularity score and 1) the features in the given metadata and 2) the detected labels in the cats & dogs photos using the Google Cloud Vision API. We assume that 'cute-looking' relates to a higher popularity score (same reasoning as in the competition). Based on findings from previous studies, we hypothesise that age, breed/colour and size influence the popularity of a pet and that less 'cute-looking' cats and dogs are found more in the less popular photos as 'cute-looking' photos receive more user clicks.

Contributions. We start by confirming former work that static features age, colour/breed and size do correlate and have impact on the popularity of a pet. We continue by showing strong positive and negative correlations of other features in photos that might influence how soon a pet is adopted. We do not prove causation; however, we do see that most features make sense and thus could still be used to try to make the pet photos more attractive.

Outline. The rest of the paper is organised as follows. Section 2 discusses the ongoing Kaggle competition and other related work. In section 3, we describe the dataset, how we obtained our own labels and how the evaluation is done. Then, we analyse and discuss our results in section 4 and 5. In section 6, we conclude and give some directions of future work.

2 RELATED WORK

To help pets find a new home, Malaysia's leading animal welfare platform 'PetFinder.my' has implemented a Cuteness meter in addition to the platform they provide to put animals up for adoption. It analyses picture composition and other factors compared to the performance of thousands of pet profiles and gives your pet's photo a score on how 'cute-looking' it is. As this Cuteness meter is still in an experimental stage, 'PetFinder.my' has created a Kaggle competition to find and implement a better algorithm that predicts the pet's popularity (Pawpularity). This Cuteness meter has the potential to help users pick a better profile picture for their pet. However, as of now, this Cuteness meter does not give any hints to the user in how to improve the score. In our research, we aim to find human-interpretable features that influence the pet's popularity score, e.g., the location/terrain, visible human body parts, pet colour and breed, toys, etc.

According to [9], the pet's popularity online is related to the length of a pet's stay in a shelter. In their study, they analysed click statistics for 892 cat adoptions and found that cats who were popular on 'PetFinder.com' (evidenced by the number of clicks per day) were also popular in the shelter environment (evidenced by the shorter length of stays). These findings make a good profile picture for a pet the more important. They also found that age and coat

¹<https://www.kaggle.com/c/petfinder-pawpularity-score>

colour were strong predictors for the number of clicks and length of stay. It seems that there exists a 'black cat' adoption bias as black cats were clicked on the least, suggesting they have a harder time of being adopted. In addition to these findings, they found some non-static features that could be manipulated by the photographer; strategic use of cat toys and photographing cats outside their cages could promote adoptions of cats who are typically overlooked.

Other research supports these findings. In [2], they used several algorithms to develop models that could predict the length of stay of cats and dogs in animal shelters. They collected over 113,000 animal records from several animal shelters around the US and saw that important indicators were age, multi-colour and large and small dog sizes. [3] conducted a smaller study on the effects of the provision of toys, cage location and static characteristics (i.e., activity level, age, sex and coat colour) on 111 cats in an animal shelter. Their analysis revealed that the cats in cages located at eye level, cats with a toy in their cage and active cats were more likely to be adopted during their 16-week study.

In the machine learning domain, research has also been done on cuteness of pets in images. In [1], they tried to improve classification of whether a cat or dog photo looks cute by employing a Gaussian Mixture Model (GMM) and using an improved Fisher Vector [8] representation. Their model was extended in a way that it could also localise the cute parts of the photo by selecting the corresponding GMM components. They found that cuteness was mostly localised at the pet's face (the eyes especially). In [6], they also tried to predict cuteness but then on a scale of 0 to 10 and with baby images instead of pets. They looked at several image features like SURF, HOG and CNNs. In this new area of research, they proposed a supervised learning scheme to find potentials of using deep learning. They did, however, not achieve high accuracies.

There is not much more work done on cuteness of images. However, there exist some more animal image studies that are related, carrying out object recognition and detection from cluttered background [7][4] and a study focusing on more fine-grained categories that discriminated between different dog breeds [5].

3 APPROACH

The dataset provided by 'PetFinder.my' includes 9912 photos of cats and dogs. Each photo has a popularity score (1-100) that is derived from each pet profile's page view statistics (duplicate clicks, crawler bot accesses and sponsored profiles excluded). It also contains binary metadata of 12 features that are devised by humans and obtained by manual labelling. These features and their descriptions² can be found in Table 1. Each feature could potentially increase or decrease the popularity of a photo. To evaluate this, we use the Pearson correlation coefficient and calculate the correlation between the popularity scores and the percentages how often a particular feature is found in bins of the photos.

The Google Cloud Vision API uses powerful machine learning models to quickly detect individual objects and faces in images. It is able to discriminate cat and dog breeds, and more. To be able to use the Google Cloud Vision API, one needs a Google Cloud Platform project with an active billing account and enable the Vision API. We set up a Google Colab Notebook and connected to the Vision

Table 1: The human-labelled metadata, how often they were labelled positive, and their Pearson correlation coefficient. All p-values are ≤ 0.05 .

Feature	Description	%	r
Accessory	Accompanying physical or digital accessory / prop (i.e. toy, digital sticker), excluding collar and leash.	7	0.182
Action	Pet in the middle of an action (e.g., jumping).	1	0.097
Near	Single pet covering over 50% of photo width or height.	86	0.019
Group	More than 1 pet in the photo.	13	0.000
Face	Decently clear face, facing front or near-front.	90	-0.002
Info	Custom-added text or labels (i.e. pet name, description).	6	-0.018
Focus	Pet stands out against uncluttered background.	3	-0.029
Collage	Digitally-retouched photo (i.e. with digital photo frame, combination of multiple photos).	5	-0.038
Eyes	Both eyes are facing front or near-front, with at least 1 eye / pupil decently clear.	77	-0.153
Blur	Noticeably out of focus or noisy, especially for the pet's eyes and face. For Blur entries, "Eyes" column is always set to 0.	7	-0.176
Human	Human in the photo.	17	-0.255
Occlusion	Specific undesirable objects blocking part of the pet (i.e. human, cage or fence).	17	-0.280

API using the private key of our Google Cloud project. Then, for each photo, we made an API call and retrieved the labels together with their certainty (as a score between 0 and 100). We put these labels and scores together in a dataframe with the image id and the popularity score, and stored this on our Google Drive. Using Google Colab's resources, the execution time was 7 hours.

After analysing the results, we found that the *label_detection* function only retrieved the top-10 labels by default. Because our analysis of cat and dog breeds did not contain a lot of samples (not a lot of breeds had enough certainty to occur in the top-10), we decided to set parameter *max_results=50* to obtain more labels (or all, if ≤ 50). We also noticed that there was a first confounder in our data regarding cats and dogs. When plotting the labels 'Cat' and 'Dog', there was a negative correlation of 0.730 for cats and a positive correlation of 0.702 for dogs. To remove this confounder, we split the data based on the 'Cat' and 'Dog' labels.

Then, we manually extracted all the cats and dog breeds that occurred more than 10 times in the data. Choosing a list of cat and dog breeds from the internet did not suffice, as the names could slightly differ. Google also did not have a list available of their labels, unfortunately.

For the correlation between the popularity score and the cat and dog breeds, we decided to bin the data into 20 bins, so scores in a range of 5 were put together. We are aware that this alters the correlations between breeds and the popularity, but we considered it fair as 100 bins (score 1-100) meant a considerably large number of empty bins (because breeds did not occur in every popularity score). We decided not to use 10 bins as this averaged too much.

²Taken from <https://www.kaggle.com/c/petfinder-pawpularity-score/data>

Table 2: Pearson correlation (with degrees of freedom N-2) between the cat & dog breeds and the popularity of the photos. All p-values are ≤ 0.05 .

Cat breed	df	r	Dog breed	df	r
British longhair	40	0.77	Small terrier	243	0.86
Ragdoll	14	0.58	Toy dog	1459	0.83
Balinese	95	0.56	Shih-poo	64	0.80
Siamese	96	0.53	Mal-shi	30	0.80
Maine coon	8	0.50	Yorkipoo	49	0.78
Birman	34	0.50	Poodle crossbr.	100	0.77
⋮	⋮	⋮	⋮	⋮	⋮
Havana brown	74	-0.36	Pinscher	61	-0.47
Korat	20	-0.47	Chihuahua	138	-0.50
Bombay	305	-0.53	Gun dog	241	-0.50
Dom. short-h. cat	5031	-0.74	Dogo guatemalt.	25	-0.52
Devon rex	21	-0.81	Working dog	2048	-0.74
Oriental shorthair	23	-0.84	Sighthound	37	-0.77

In addition to the cat and dog breeds, we also computed two full lists of correlations (cats/dogs data) from all the other features the Vision API had found. Labels found in less than 10 images were again excluded. In our analysis, we looked at the features with the strongest positive and negative correlations. We also selected some other labels, e.g., ‘Toy’ and ‘Human body’, based on features of previous work and the metadata of the ‘PetFinder.my’ dataset.

Full code is made available in our Github repository³. Here, one can also find bar plots for each of the features in this paper, showing the correlation between this feature and its popularity score.

4 RESULTS & ANALYSIS

4.1 PetFinder.my metadata features

Table 1 shows the 12 features from the human-labelling and their correlations to the pet popularity. The middle columns give a specification how the features were labelled and how often they were labelled as ‘1’. As can be seen, some features were found very often or very little in the pet photos. This indicates that the distributions are quite unbalanced. Most pets were near the camera with their face pointed towards it. In the end, this did not influence popularity, as many pets who did not have this feature also obtained high popularity scores. Actually, all correlations are rather weak. When plotting the relation between a feature and the popularity, one can see features are found at all ranges of scores. We suspect there are other unidentified features in the photos that influence all correlations, therefore giving unexpected results. The highest correlations are for ‘Occlusion’ (-0.280) and ‘Human’ (-0.255). The feature ‘Eyes’ gives surprising results as this involves a negative correlation. One would expect the feature to influence people’s opinions on cuteness positively. It also contradicts [1] where they localised cuteness to be especially around the eyes.

³<https://github.com/.../TxMM-PawpularityContest>

4.2 Google Vision API features

After extracting the Google Vision labels, it became clear that there was a strong negative correlation (-0.730) for cats and a strong positive relation (0.702) for dogs. We can conclude from this that dogs are in general more popular on ‘PetFinder.my’ than cats as they relatively obtain higher Pawpularity scores. We removed this confounder as users of the Cuteness meter probably do not wish to see the hint ‘change your pet to a dog to obtain a better score’.

Table 2 shows the correlations between a breed and the popularity score after splitting the data. Breeds with a p-value higher than 0.05 were removed from the list, e.g., Terrier had the highest correlation ($r(4445) = 0.91$, $p = 0.50$), but this could have occurred just by random chance. Looking at the remaining breeds, an order of breed popularity can be derived from the table. In general, breeds had a positive correlation on the pet popularity. We think this is due to a clearer image and the Vision API being better able to recognise breeds in higher popularity scores. We also see small dogs appear relatively higher than big dogs, confirming previous work that small dogs are more easily adopted.

What we see from Table 2 is that there is a significant preference for breed when deciding to click on a pet’s profile page. In general, small fluffy dogs as the Shih-poo are found much more often in the popular dog images, and so are Sighthounds and Dogo Guatemaltecos less popular breeds. When looking at the posture of these two, it is understandable that these breeds receive less clicks from users. For cats, negative correlations are found for short-haired cats and black cats (fulfilling the ‘black cat’ adoption bias). The Devon Rex is a cat with remarkable large ears and slender bodies (apparently not found very attractive). Also, Bombay cats with their black fur

Table 3: Highest and lowest Pearson correlations (with degrees of freedom N-2) between features and their popularity scores. In the middle, some features were selected to compare related work. All p-values are ≤ 0.05 .

Cat features	df	r	Dog features	df	r
Plush	22	0.52	Smile	88	0.65
Stuffed toy	50	0.46	Puppy	1551	0.60
Bed	141	0.45	Close-up	10	0.57
Wood	1422	0.44	Puppy love	335	0.56
Toy	163	0.44	Finger	37	-0.03
Human leg	378	-0.34	Hand	51	-0.07
Hand	140	-0.36	Toy	166	-0.11
Human body	62	-0.48	Human leg	162	-0.38
Finger	134	-0.55	Foot	354	-0.57
Net	28	-0.61	Soil	258	-0.66
Kennel	126	-0.61	Asphalt	217	-0.67
Road	75	-0.65	Shrub	25	-0.69
Soil	91	-0.67	Full-size car	13	-0.70
Darkness	45	-0.67	Plant	874	-0.71
Night	8	-0.67	Rectangle	249	-0.74
Sidewalk	83	-0.70	Landscape	55	-0.75
Newspaper	27	-0.71	Dog collar	1494	-0.75

are not in luck when it comes to popularity. The more popular cat breeds are again fluffy and more long-haired.

In Table 3, one can see the strongest correlations of the other features we examined. The ones in the middle were selected based on features of previous studies. We see that the finding that cats with toys are more likeable is reproduced. For dogs, it does not seem to matter whether there is a toy in the photo or not. Furthermore, we see a stronger negative correlation in this experiment whether a human is in the photo or not. This finding is more reliable than when we looked at the label 'Human' in the metadata as we will discuss in Section 5. This relation does not seem to hold for hands and fingers in dog photos, though.

What can be derived from these features, are some hints on how to improve pets' profile pictures. For cats, it is best to surround them with softness (plush or a bed) or on a wooden floor when it comes to making their profile pictures. Cats are not that popular when they are on the street, on soil or on a newspaper. Also, the Vision API sometimes misinterpreted 'Net' for bars of a cage. So, a photo of a cat in a kennel or cage is also not beneficial. For dogs, we see a similar pattern that soil and asphalt decrease the popularity of a dog. We also see features like 'Shrub', 'Plant' and 'Landscape' having a negative correlation. Important when photographing dogs is that they are made close-up and do not contain much landscape around them. Dog photos were also more popular when the Vision API could recognise a smile somewhere in the photo.

In addition, these findings suggest that age is also an important indicator for adoption. Photos of puppies were much more popular due to its high correlation. Another interesting finding is the strong negative correlation of 'Dog collar'. This might be because it could indicate the dog already has an owner and does not need one, or because puppies often do not have a collar yet.

5 DISCUSSION

Our findings support previous work that popularity (and the adoption) of a pet quite depends on its age, breed/colour and size. Fortunately, we also found some features that are manipulable by the photographer, such as the surroundings and the flooring. The use and fairness of a cuteness ranker is debatable, because there are a lot of static features. Is it still fair to make a cuteness ranker in which users insert a photo and try to improve their pet's Pawpularity score, when that score is not only based on camera quality features, pet positions and attributes/objects, but also on something unchangeable like its age, breed/colour and size?

These uncontrollable features need to be considered when designing an algorithm that ranks pets on cuteness. It is also important to inform the users of the tool about the features that are used to determine the score, and make them aware that, for example, a perfect score for their black pet does not necessarily lead to instant adoption.

We would also like to discuss some limitations of our research. Firstly, we studied cuteness on cats and dogs by looking at their popularity score. It is reasonable to assume that the cuteness of an animal relates to their popularity score, because of the previous works that show cute photos attract people's attention more, receive more user-clicks and have a bigger chance of being adopted. However, this relationship between cuteness and the popularity




Figure 1: (a&b) Least and most popular 'Newspaper' cat photo. (c) Random photo with score 100. Visible cages and certain flooring, such as newspapers, decrease the attractiveness to humans.

score is never perfectly aligned and thus debatable. Our research shows that dogs were more popular than cats on 'PetFinder.my', but this does not necessarily mean that they are also considered more 'cute' than cats.

Another limitation involves the datasets we have used. The 'PetFinder.my' dataset was created for the same goal: trying to improve the cuteness ranker. We, however, shifted our approach to a descriptive mining task instead of building a prediction model. Additionally, a key factor that is missing (and becomes more important in our research) is any information about the procedure of the human labelling. 'PetFinder.my' did not indicate how they obtained their labels. That makes the first part of our work, the analysis on the metadata, less trustworthy. If we would have had the time and resources, we could have manually labelled the data ourselves and ensure validity. This was, however, unfeasible in the scope of this project. Lastly, we used the Google Vision API to label our photos, which comes with its own limitations, as the API is not perfect and sometimes mislabels objects. It is, however, one of the state-of-the-art tools.

6 CONCLUSION & FUTURE WORK

In this research, we took a Multimedia mining approach to find features that could influence pets' profile picture popularity on online adoption platforms. We found that static features such as age, breed/colour and size correlate with the cat and dog popularities. Furthermore, we discovered which breeds are positively correlated and which ones negatively. Especially, small fluffy dogs and long-haired cats had high popularity scores. Breeds that were less popular were black, short-haired, or unusually shaped animals (such as the Devon Rex, Dogo Guatemalteco or Sighthound) which could thus be considered 'less' cute-looking. We also found several other features in photos that show negative correlations, such as the flooring (soil, asphalt, newspaper) or landscape in the photo (not being close-up). This knowledge was obtained by extracting our own labels with the Google Vision API. The correlations from the first part were more surprising, we suspect them to be less trustworthy.

Since we did not have information about this labelling process, future research could do the human-labelling of this metadata again to ensure validity. The cuteness of pets could also be studied from a different angle, for example, in the field of deep learning where new features are extracted, finding improvements in the localisation of cuteness in photos[1], or carrying out user studies by asking

people directly what determines cuteness in cat and dog photos for them. Lastly, future work could try to improve the current cuteness ranker by implementing a new machine learning model.

REFERENCES

- [1] Y. Bao, J. Yang, L. Cao, H. Li, and J. Tang. 2014. Cuteness Recognition and Localization in the Photos of Animals. In *Proceedings of the 22nd ACM International Conference on Multimedia*. 1193–1196.
- [2] J. Bradley and S. Rajendran. 2021. Increasing adoption rates at animal shelters: a two-phase approach to predict length of stay and optimal shelter allocation. *BMC Veterinary Research* 17, 70 (2021).
- [3] J.M. Fantuzzi, K.A. Miller, and E. Weiss. 2010. Factors relevant to adoption of cats in an animal shelter. *Journal of Applied Animal Welfare Science: JAAWS* 13, 2 (2010).
- [4] F. Fleuret and D. Geman. 2008. Stationary Features and Cat Detection. *Journal of Machine Learning Research* 9 (2008), 2549–2578.
- [5] A. Khosla, N. Jayadevaprakash, B. Yao, and L. Fei-Fei. 2012. *Stanford Dogs Dataset*. <http://vision.stanford.edu/aditya86/ImageNetDogs/>.
- [6] P. Makula, A. Kumar, and S. Mukherjee. 2018. Measuring Level of Cuteness of Baby Images: A Supervised Learning Scheme. In *Multimed Tools Appl*, Vol. 77. 16867–16885.
- [7] O.M. Parkhi, A. Vedaldi, C.V. Jawahar, and A. Zisserman. 2011. The truth about cats and dogs. In *2011 International Conference on Computer Vision*. 1427–1434.
- [8] F. Perronnin, J. Sánchez, and T. Mensink. 2010. Improving the fisher kernel for large-scale image classification. In *ECCV*.
- [9] M.K. Workman and C.L. Hoffman. 2015. An Evaluation of the Role the Internet Site Petfinder Plays in Cat Adoptions. *Journal of Applied Animal Welfare Science* 18, 4 (2015), 388–397.