Leaf N' Meow: Predicting Plant Toxicity to Cats

Deanna Fedash
dfedash@ucsc.edu
University of California Santa Cruz
Santa Cruz, California, USA

David Abramov
dabramov@ucsc.edu
University of California Santa Cruz
Santa Cruz, California, USA

Figure 1: *Monstera deliciosa* makes for an aesthetic houseplant, but contains calcium oxalates, which are toxic to cats.

ABSTRACT

Unknown to many cat owners, some common houseplants are lethally toxic to their fuzzy companions. We examine several hundred plant species known to be either toxic or non-toxic to cats, along with a number of different plant traits and characteristics. After downloading, curating and cleaning our source data, we trained and tested several machine learning models in order to predict whether or not a plant is toxic to cats – which is much more humane than testing toxicity directly. We found that the decision tree algorithm had the highest performance metrics out of our experimental models. Hopefully this work can help inform other researchers in the scientific community as to what measurements to better collect across species in order to make our model more predictive.

ACM Reference Format:
Deanna Fedash and David Abramov. . Leaf N' Meow: Predicting Plant Toxicity to Cats. In *Proceedings of ACM Conference (Conference'17).* ACM, New York, NY, USA, 7 pages.

© 0
ACM ISBN 978-x-xxxx-xxxx-x/YY/MM…$15.00

1 INTRODUCTION

Despite being obligate carnivores and partially domesticated, house cats can exhibit some undesirable behaviors such as munching on plants and sometimes using them as their bathroom in addition to their litter box [17]. Unfortunately some plants that are trendy to decorate with are not always friendly to cats. The *Monstera deliciosa* leaves shown behind the cats in **Figure 1** contains insoluble calcium oxalate crystals (CaOx) that can irritate cat oral tissues, cause inflammation in and around the mouth, as well as induce vomiting [1]. CaOx is accumulated across a diverse range of photosynthetic organisms ranging from algaes to big leafy trees. These crystals are generally embedded within plant cell walls or are contained within vacuoles. It is presumed that some species have evolved an increased expression of CaOx accumulation as a defense mechanism in order to deter herbivores. In addition to the negative impacts to grazing animals, CaOx crystals can be present in human diets and lead to kidney stones [9].

Fortunately not all plants are toxic to cats, and some even have a positive symbiotic relationship with cats. Catnip, *Nepeta cataria*, is often given to cats as a treat and causes them to display some interesting behaviors. Catnip and other mint related plants contain semiochemicals that attract and stimulate a cat's strong olfactory senses. Evidence has shown that the allomones produced by these herbs simultaneously repel insects as well as attract predators such

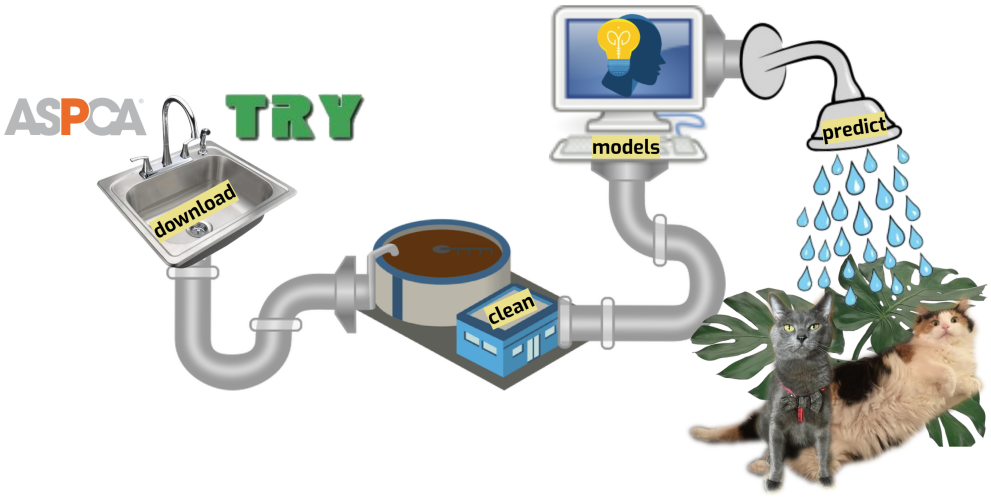

Figure 2: Illustration of our data processing pipeline. We first downloaded the list of (non-)toxic plants on the ASPCA website, and for each of the species we found the corresponding ID in the TRY plant trait database. With those plant IDs, we were able to download plant traits that we thought could be predictive. Next, we cleaned the data by identifying and removing duplicate plants in the ASPCA list, converting categorical attributes to numerical equivalents, imputing missing values, and normalizing the final values. We then trained five different machine learning models with these cleaned and organized data: Logistic Regression, Support Vector Machines, K-Nearest Neighbors, Decision Trees, and Gradient Boosted Decision Trees. Finally, we ran these models with test data to measure the reliance of their predictability.

as the cat, who might then eat the pests who would otherwise snack on the plants. [4].

1.1 Motivation

An often unfortunate overlooked detail in pet ownership is that certain ornamental plants that decorate people's homes and gardens are toxic to their beloved creatures. We want to be able to easily look up a plant and determine whether or not it is toxic or non-toxic to cats. Furthermore, we would like to be able to identify any predictive plant attributes if it is otherwise unknown whether or not a plant is toxic to cats.

People who are spontaneous plant shoppers and also have cats, such as ourselves, often do not have enough self control to check whether or not a plant is safe for our cats before we buy them. Later when we get home, we have to somehow separate our cats from our plants, which is not always convenient or match the ideal location in our home that we had in mind. We would like to help other cat owners out there make the most responsible decisions for their pets and to prevent any unnecessary harm and costly trips to the veterinarian.

We hope that we can predict to a certain extent whether a plant is toxic or non-toxic to cats. We expect to find at least a few plant traits that can be used as a predictive measure for plant toxicity to cats, as there are possibly more than one biological pathway that can make a plant toxic. A good model would minimize false negatives (specificity) and avoid results labeled as non-toxic when really they are toxic, especially for plants that have an unidentified toxicity to cats. Based on the nature of our data, we expect to be able to identify toxic and non-toxic plants based on their common or scientific name if they are included in the original dataset. However, for other plants that are not included in the toxicity data, we hope to identify some plant traits that can be used to predict toxicity.

1.2 Related Work

There are quite a number of applications available on Android and iOS that predict plant names based on pictures. Here are some we looked into: PlantNet[11][5][12], LeafSnap[15] and PlantSnap. These all use "machine learning technology," which likely means a pre-trained image classification network tuned for plant species such as LeafNet [3]. We are not very interested in using image recognition in this project, but it would be a beneficial feature as PlantSnap has been documented in the medical literature for being used to diagnose plant-related illnesses[7][14]. While useful for identification, there is not an app that directly connects plant identification with cat toxicity and potential medical reactions.

Scientific studies have been done on specific plants and their toxicity to cats. Lilies are notoriously deadly, as only two leaves can kill a cat, damaging the liver. The exact mechanism is not well known [10][8][16]. An overview of potentially life-threatening poisonous plants in dogs and cats provides information about some of the most common poisonous plants, with geographic distribution, toxicity, clinical signs, and treatment [18]. The paper Household Food Items Toxic to Dogs and Cats goes over a number of different foods we find tasty but cats and dogs find deadly, such as chocolate, grapes, onions, garlic and alcohol, but they mostly focus on dogs [6].

We believe our work will be novel because we have not been able to find any examples in the literature or in an app store that accomplishes our task. Our current method of looking up a plant's toxicity to cats is with a quick Google search, which is not always successful. In addition, since not much is known about some of the mechanisms that make plants toxic, our results could help researchers better understand and uncover biological pathways to target for treatment, identify new compounds, or alter plant genetics to erase toxicity.

2 ANALYSIS PIPELINE

There were a number of data mining methods we employed in order to build, train and test machine learning models to predict a plant's toxicity to cats. Our data pipeline is illustrated in **Figure 2** and is broken down into four steps: download, clean, models, and predict.

2.1 Data Sources

On their website, the American Society for the Prevention of Cruelty to Animals (ASPCA) provides a list of plants that are toxic and non-toxic to cats [2]. The ASPCA shares around 1,000 plant species with their common and scientific names as well as familial taxonomy. For each species they also provide an image, and a description that includes clinical signs and possible toxic compounds in the plant.

We accessed data in the TRY Plant Traits Database in addition to the toxicity labels from the ASPCA database. The TRY database contains nearly every plant species known to science, covering at least 100,000+ plant species, with over 2,000 measured traits including morphological, biochemical, and physiological features[13]. Around 400 attributes have a sample size greater than 1,000 and range from things like plant vegetative height, leaf phenology, mycorrhiza type, to human usage types. The data was made available to download after registering an account.

2.2 Cleaning and Wrangling

We used several data mining techniques in order to wrangle our data into a clean, manageable number of attributes and species. First, we looked for overlapping species in the ASPCA and TRY databases, and from there began to narrow down the attributes to select, as to not later fall prey to the curse of dimensionality. While cleaning, we found that there were many duplicates in the ASPCA list due to misspelling, narrowing down our number of labeled species to around 500 plants.

From the full TRY Trait Database, we selected 53 numerical and categorical attributes that we thought might be predictive. Since not each trait completely covered the list of labeled species from the ASPCA, we had to impute missing values, using the mean for numerical attributes and mode for categorical. Mean and mode were also used to combine duplicate measurements for the features. We performed logistic regression on the selected attributes individually to determine which ones to remove to further reduce the dimensionality and determine the most predictive traits. We measured the attributes against the amount of the labeled species in the ASPCA they covered and their accuracy score from the logistic regression model. We compared the performance of the selected features with the performance of the full 53 features on each of our models.

2.3 Machine Learning Models

After cleaning our data and narrowing down which traits to use, we trained and tested five machine learning models: logistic regression, support vector machines, k-nearest neighbor, decision trees, and gradient boosted decision trees. We chose these models for different reasons. All models are able to handle two-class classification and provide easily interpretable results. While we would expect neural networks to perform pretty well with these data, it would be hard to evaluate what exactly the model is learning. Since modeling algorithms each have their individual strengths and weaknesses, we were interested in comparing the performance of linear models, such as support vector machines and logistic regression, and non-linear models like decision trees and k-nearest neighbor. In addition, each of these models have readily available software support. The k-nearest neighbor algorithm has been shown to be able to handle noisy data like ours. A hierarchical method like decision trees seems well suited for this type our type analysis.

2.4 Make Predictions

We verified that the model was good by training it on a subset of our cleaned data, and then ran it again with the test subset. We used an 80:20 split for our training and test data and split the data multiple ways following the k-fold cross validation method. We did this in order to avoid the models over-fitting to the training data. We verified the accuracy and specificity of our model using plants we know the cat toxicity of but were not used to train the models. We also compared the relative performance for each model using an ROC curve along with their true negative rate (specificity). We chose to maximize the true negative rate to avoid classifying a toxic plant as non-toxic.

3 ANALYSIS

After cleaning our dataset we trained and tested several machine learning models on the plant traits we accessed from the TRY database along with the toxicity label from the ASPCA website. We first ran a logistic regression on all of the traits that we extracted from the TRY data to compare their relative predictability. In order to generate more robust models, avoid overfitting, and prevent the curse of dimensionality, we trained our models again using only a subset of the most predictive traits. We compared models across a few different performance metrics.

3.1 Results

The results from the logistic regression analysis on the individual categorical and numerical features are show in **Table 1** and **Table 2** respectively. The top predictive features for toxicity were Root Calcium (74.7%), Latitude (69.1%), Longitude (71.0%), Hemeroby, (62.6%), Urbanity (62.7%), and Post Fire Seedlings (78.8%).

We trained and tested the data on five different machine learning models: logistic regression, support vector machines, k-nearest neighbors, decision trees, and a gradient boosted decision trees. Results for both the full dataset and the selected feature subset are shown in **Table 3**. Models ran with the subset showed an increase in both average accuracy and AUC score show across all models. When comparing specificity between models trained on all features and just the predictive subset, logistic regression and decision trees

Feature	Average Accuracy (%)	Number of Data Points
Exposure to Freezing in Natural Range	44.9	223
Hemeroby Categories	62.6	206
Leaf Type	56.7	2246
Native Continent	48.4	124
Order	51.4	111
Phylogenetic Group	44.8	484
Shade Tolerance	47.3	296
Urbanity	64.4	87
Family	62.7	528
CaCO3 Tolerance	51.3	76
Anaerobic Tolerance	45.1	78
Drought Tolerance	55.2	87
Fire Resistance	46.7	193
Hedge Tolerance	37.2	78
Livestock Bloat	60.0	57
Mycorrhiza	61.6	86
Palatable to Browse Animal	46.0	63
Palatable to Graze Animal	42.3	57
Palatable to Human	38.8	80
Post Fire Seedlings	78.8	78
Salinity Tolerance	59.7	79
Toxicity from TRY	54.9	101

Table 1: Performance metrics across categorical features. Average accuracy was recorded from running logistic regression individually for each attribute with k-cross folds validation. Number of data points for each feature is also shown.

showed an increase, whereas the support vector machines and k-nearest neighbors models showed a decrease. ROC plots comparing the subsets and all features for logistic regression, support vector machines, and k-nearest neighbors are shown in **Figure 3**.

We found that decision trees performed the best with the highest average accuracy from k-folds crossover validation, the highest true negative rate (specificity), as well as the greatest area under the curve (AUC) score from the ROC. The Boosted Decision Tree model when run on the subset, showed a small increase in average accuracy compared to Decision Tree run on the subset, however there was a decrease in both specificity and AUC score. ROC plots comparing the Boosted Decision Tree on the subset and Normal Decision Tree on the full features are show in **Figure 4**. The AUC score between the models are the same, but the boosted model had a more curvilinear appearance.

3.2 Interpretation

When the models were run on the subset of more predictive features, there was a increase in performance in the metrics of accuracy and the AUC score. The logistic regression and k-nearest neighbor models run on the full features dataset had a higher true negative

Feature	Average Accuracy (%)	Number of Data Points
Altitude	51.8	1113
Annual Moisture Balance	37.5	77
Annual Precipitation	57.2	1121
Annual Radiation	59.6	77
EW Index	30.0	30.0
Forest Productivity Class	48.7	192
Frost Tolerance (Days)	48.2	80
Frost Tolerance (Celsius)	37.2	78
GDD0	48.2	46
Growth Rate	60.9	74
Growth Season Precipitation	40.4	77
Growth Season Temperature)	48.5	130
Growth Season VPD	46.9	130
Habitat (Dry)	54.4	72
Habitat (Medium)	61.2	72
Habitat (Moist)	61.2	72
Hemeroby Numerical	61.2	85
Latitude	69.1	5150
Leaf Area Index	55.8	87
Leaf Calcium	12.9	170
Leaf Width	61.9	707
Life Span (years)	41.4	329
Longitude	71.0	5153
Mean Annual Temperature	44.0	784
Number of Floristic Zones	54.6	83
Number of Native European Countries	39.4	64
Potential Evapotranspiration	59.4	113
Root Calcium	74.7	169
Shoot Calcium	55.7	32
VPD	45.6	77

Table 2: Performance metrics across numerical features Average accuracy from running individual features on a logistic regression model to predict toxicity. Features were also evaluated on the number of data points they covered.

rate. We believe this could be explained by some features that were not predictive on their own but when combined with others, increased the specificity of our models.

The non-boosted decision tree algorithm had the highest performance metrics out of the models we trained and evaluated, not including the average accuracy boosted decision tree. We believe

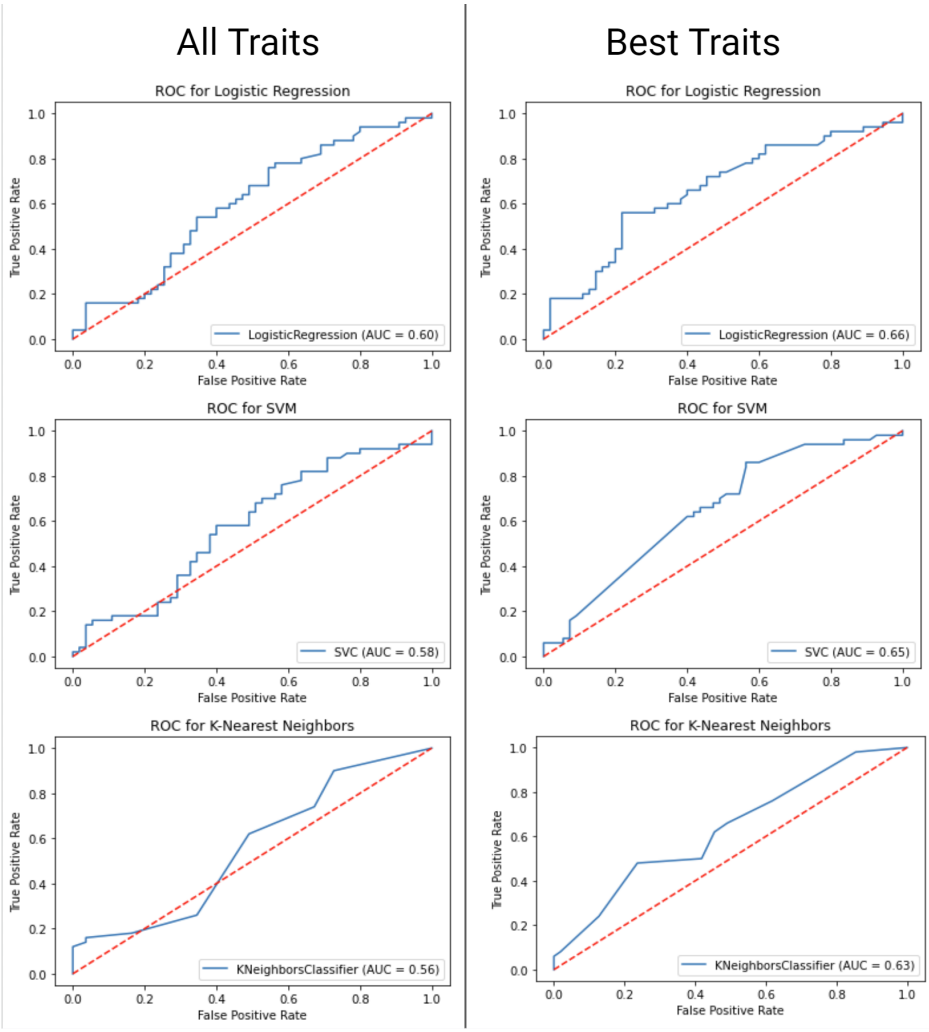

Figure 3: ROC plots for logistic regression, support vector machines and k-nearest neighbors before and after selecting top traits. Models trained with all of the traits in **Figures 2 and 3** are shown on the left, and models trained with only the most predictive are shown on the right. Models trained on only the best traits had a boost of about 0.06-0.07 on the AUC score.

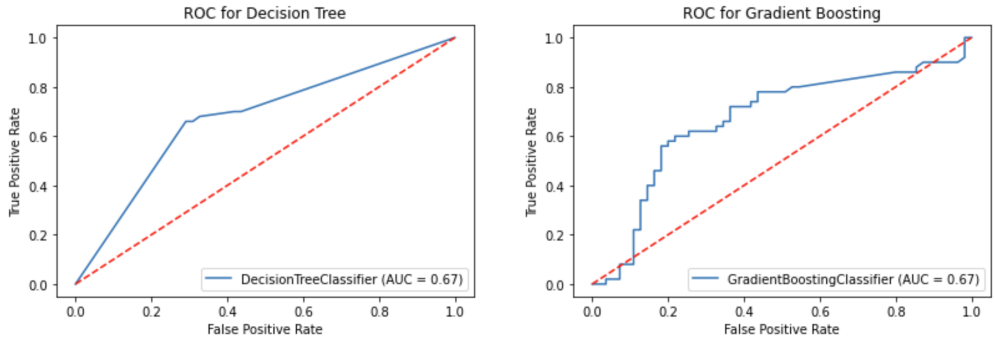

Figure 4: ROC plots for Decision Tree models. On the left is the ROC for the vanilla decision tree trained and tested on all of the features in **Tables 2 and 3.** On the right is the ROC for the decision tree with gradient boosting trained on only the most predictive plant traits. Both models had the same AUC score of 0.67, however the boosted model appears slightly higher resolution and more curvilinear.

that this is partially due to the lack of linear separability between the plants labeled as toxic versus non-toxic as indicated by PCA analysis of the clean dataset. Based on the results of all the models, we would need to calculate a likelihood value for plant toxicity in order to extend it to species outside of the ASPCA database.

The average accuracy of root calcium was expected based on the scientific evidence behind lily toxicity, as well as calcium oxalate as it is mentioned as a plant toxin in the ASPCA database. Latitude and longitude were also predictive, indicating that the species range of a particular plant and ancestral cats could have had a evolutionary relationship. In addition, hemeroby and urbanity, measures of human influence on the environment and the adaptiveness of plants to human centers were predictive. We believe that there is a measured effect of the partial domestication of cats over the past several thousand years.

We expected family would be more predictive of plant toxicity, however, after looking at the data as a pivot table, we noticed that some families were completely toxic or non-toxic, while others had a more even distribution. Post fire seedlings, or whether plant sprouts emerges after a wildfire, was surprising to us in its accuracy of 78.8%, but a caveat is that we did not have a ton of data points in this category.

4 CONCLUSION

Predicting plant toxicity to cats based on plant traits is not an easy problem to solve. There are many different biological pathways, known and unknown, that can lead a particular plant to be toxic to cats – with calcium being a relatively good predictor. In the context of evolution, it is interesting to note that some plants have a symbiotic relationship with cats, such as the insect-repelling and cat-enticing aroma of catnip, whereas other plants have instead developed defense mechanisms. There appears to be a relationship between geographical coverage of a plant species (habitat) and its toxicity to cats – perhaps plants whose range overlaps with ancestral cats selectively evolved grazing defense mechanisms, or vice versa, cats became adapted to the plants they encountered (survivorship bias). Considering that cats have been partially domesticated in the presence of humans, it is interesting to note that human-flora relationships have a predictive affect on plant toxicity to cats.

We cleaned and processed around 50 different plant traits for about 500 plants we knew were toxic and nontoxic to cats, and trained several different machine learning models. Linear models such as logistic regression were limited due to the data being non-linearly separated and the model is more noise sensitive. The non-boosted decision tree had the best performance, however, if we were to use this model to predict on plant species that were not in the ASPCA database, we would want to calculate a likelihood value for toxicity. We were able to find a few predictive traits for toxicity from our based on the results of the individual feature logistic regression analysis.

4.1 Future Work

While our decision trees performed pretty well, based on our results we believe that more work needs to be done in order to robustly predict toxicity to cats across the plant kingdom. Part of the issue with this project was the limited number of plants labeled with toxicity in the ASPCA database, and then the lack of uniform coverage of corresponding traits from the TRY database. Plant traits in the TRY database was sparse for many plants, and surprisingly messy within a particular attribute. Data points tended to be heavily skewed towards some plants such as the maple tree, perhaps due to its economic viability. In addition, there were many other plant traits within the TRY database that we could look into that we did not initially consider. The scientific community could use this work as motivation to collect calcium measurements within and across more plant species as it was one of our most predictive attributes. There could also be other potential toxic compounds within TRY that we did not consider. A more sophisticated project would be to perform some genetic analysis in order to identify specific proteins or toxins that are produced as plant defense mechanisms.

Model	Average Accuracy (%)	Specificity	AUC Score
Logistic Regression (All Features)	58.3	67.3	0.60
Logistic Regression (Subset)	61.9	69.1	0.66
Support Vector Machines (All Features)	60.9	74.5	0.58
Support Vector Machines (Subset)	64.0	56.3	0.65
K-Nearest Neighbors (All Features)	61.4	96.4	0.58
K-Nearest Neighbors (Subset)	64.2	76.4	0.63
Decision Trees (All Features)	63.8	69.1	0.67
Decision Trees (Subset)	65.4	81.8	0.70
Decision Trees (Subset, Boosted)	67.1	78.2	0.67

Table 3: Performance metrics captured across tested models: logistic regression, support vector machines, k-nearest neighbors, and decision trees. Each model was initially trained on all of our features (as shown in Table 2) and again on only a subset of the most predictive variables from the logistic regression, having an accuracy above our threshold of 50%. In addition, for decision trees we included a boosted model as well, which had the greatest accuracy (67.1%). However, our vanilla decision tree trained on the most predictive subset had a higher specificity of 81.8% as well as a 0.70 AUC score, which was higher.

REFERENCES

[1] [n.d.]. Ceriman. https://www.aspca.org/pet-care/animal-poison-control/toxic-and-non-toxic-plants/ceriman

[2] [n.d.]. Toxic and Non-Toxic Plant List - Cats. https://www.aspca.org/pet-care/animal-poison-control/cats-plant-list

[3] Pierre Barré, Ben C Stöver, Kai F Müller, and Volker Steinhage. 2017. LeafNet: A computer vision system for automatic plant species identification. *Ecological Informatics* 40 (2017), 50–56.

[4] Sebastiaan Bol, Jana Caspers, Lauren Buckingham, Gail Denise Anderson-Shelton, Carrie Ridgway, CA Tony Buffington, Stefan Schulz, and Evelien M Bunnik. 2017. Responsiveness of cats (Felidae) to silver vine (Actinidia polygama), Tatarian honeysuckle (Lonicera tatarica), valerian (Valeriana officinalis) and catnip (Nepeta cataria). *BMC veterinary research* 13, 1 (2017), 70.

[5] Pierre Bonnet, Alexis Joly, Hervé Goëau, Julien Champ, Christel Vignau, Jean-François Molino, Daniel Barthélémy, and Nozha Boujemaa. 2016. Plant identification: man vs. machine. *Multimedia Tools and Applications* 75, 3 (2016), 1647–1665.

[6] Cristina Cortinovis and Francesca Caloni. 2016. Household food items toxic to dogs and cats. *Frontiers in veterinary science* 3 (2016), 26.

[7] Mehmet Esat Ferhatlar, Vahide Aslankoç, and Neşe Kaya. [n.d.]. Use of Smartphone Applications in Herbal Poisoning. ([n. d.]).

[8] Kevin T Fitzgerald. 2010. Lily toxicity in the cat. *Topics in companion animal medicine* 25, 4 (2010), 213–217.

[9] Vincent R Franceschi and Paul A Nakata. 2005. Calcium oxalate in plants: formation and function. *Annu. Rev. Plant Biol.* 56 (2005), 41–71.

[10] Jeffery O Hall. 2013. Lilies. In *Small animal toxicology*. Elsevier, 617–620.

[11] Alexis Joly, Pierre Bonnet, Hervé Goëau, Julien Barbe, Souheil Selmi, Julien Champ, Samuel Dufour-Kowalski, Antoine Affouard, Jennifer Carré, Jean-François Molino, et al. 2016. A look inside the Pl@ ntNet experience. *Multimedia Systems* 22, 6 (2016), 751–766.

[12] Alexis Joly, Hervé Goëau, Pierre Bonnet, Vera Bakić, Julien Barbe, Souheil Selmi, Itheri Yahiaoui, Jennifer Carré, Elise Mouysset, Jean-François Molino, et al. 2014. Interactive plant identification based on social image data. *Ecological Informatics* 23 (2014), 22–34.

[13] Jens Kattge, Sandra Diaz, Sandra Lavorel, I Colin Prentice, Paul Leadley, Gerhard Bönisch, Eric Garnier, Mark Westoby, Peter B Reich, Ian J Wright, et al. 2011. TRY–a global database of plant traits. *Global change biology* 17, 9 (2011), 2905–2935.

[14] Breelan M Kear, Richard W Lee, Sanford B Church, Fady A Youssef, and Anthony Arguija. 2020. Ingestion of A Common Plant's Leaves Leads to Acute Respiratory Arrest and Paralysis: A Case Report. *Clinical practice and cases in emergency medicine* 4, 3 (2020), 371.

[15] Neeraj Kumar, Peter N Belhumeur, Arijit Biswas, David W Jacobs, W John Kress, Ida C Lopez, and João VB Soares. 2012. Leafsnap: A computer vision system for automatic plant species identification. In *European conference on computer vision*. Springer, 502–516.

[16] Cathy E Langston. 2002. Acute renal failure caused by lily ingestion in six cats. *Journal of the American Veterinary Medical Association* 220, 1 (2002), 49–52.

[17] Jacqueline M Ley. 2015. Normal but unwanted behavior in cats. *Feline Behavioral Health and Welfare* (2015), 319.

[18] Lynn M Milewski and Safdar A Khan. 2006. An overview of potentially life-threatening poisonous plants in dogs and cats. *Journal of Veterinary Emergency and Critical Care* 16, 1 (2006), 25–33.