JOURNAL OF
EMERGING INVESTIGATORS

Article

Predicting college retention rates from Google Street View images of campuses

Rushil Dileep[1], Mithil Patil[2], Tarun Mahesh[3], Atharv Deshmukh[4], Samuel Liu[2], Aarush Narayanan[5], Larry McMahan[6]

[1] Irvington High School, Fremont, California
[2] BASIS Independent Silicon Valley, San Jose, California
[3] Hillsborough High School, Hillsborough Township, New Jersey
[4] BASIS Independent Fremont, Fremont, California
[5] Mountain House High School, Mountain House, California
[6] Computer Science & Engineering Department, Aspiring Scholars Directed Research Program, Fremont, California

SUMMARY

Every year, around 40% of undergraduate students in the United States discontinue their studies, resulting in a loss of valuable education for students and a loss of money for colleges. Even so, colleges across the nation struggle to discover the underlying causes of these high dropout rates. In our paper, we discuss the use of machine learning to find correlations between the built environment factors and the retention rates of colleges. We hypothesized that one way for colleges to improve their retention rates could be to improve the physical characteristics of their campus to be more pleasing. We used image classification techniques to look at images of colleges and correlate certain features like colors, cars, and people to higher or lower retention rates. With three possible options of high, medium, and low retention rates, the probability that our models reached the right conclusion if they simply chose randomly was 33%. After finding that this 33%, or 0.33 mark, always fell outside of our 99% confidence intervals built around our models' accuracies, we concluded that our machine learning techniques can be used to find correlations between certain environmental factors and retention rates.

INTRODUCTION

In the pursuit of higher education, the importance of college cannot be overstated. It serves as a transformative period in a young individual's life, equipping them with knowledge, skills, and experiences that lay the foundation for future success. Consequently, understanding the factors that contribute to college students' ability to persist and succeed in their academic endeavors becomes crucial.

The Department of Education defines retention rate as the proportion of undergraduate students who continue their college education the following year (1). These rates provide valuable insights into the effectiveness of educational institutions in supporting their students throughout their academic journey. By examining retention rates, educators, policymakers, and school administrators can identify areas of improvement, implement targeted interventions, and ultimately foster an environment conducive to student success. In the fall of 2021, United States post-high school institutions had a retention rate of 75.6% (2). This highlights the pressing need to understand and address the factors that lead to almost a quarter of first-year undergraduate students leaving universities.

Numerous studies have been conducted to analyze the factors that contribute to this low retention rate. Studies from the American College Testing (ACT) Incorporated show that students with poor academic attributes, such as low high school grade point averages, low ACT scores, or even low academic confidence, were more prone to discontinue their college education (3). Furthermore, studies in the Research in Higher Education Journal indicated that students in certain demographics, such as Hispanic, or with negative high school experiences, such as a high number of suspensions, are more inclined to drop out of college (4). These findings highlight that both academic and nonacademic attributes of students contribute to the issue of low retention rates in higher education.

Existing research has shed important light on various factors influencing retention rates. Many colleges have established programs that help students who need academic assistance, mental health and emotional assistance, or financial assistance, and studies show that these programs are very beneficial for guiding a student who needs aid towards a successful future (5). However, these numerous studies focused heavily on student-centric factors rather than factors related to the college itself. Our study aims to explore an overlooked, yet potentially significant factor: the built environment of the college campus (3).

In our study, we defined the built environment as the general visual appearance of the college campus. We hypothesized that more pleasing physical characteristics on a college campus leads to higher student retention rates. The rationale behind our hypothesis lies in the understanding that the built environment can have a profound impact on individuals' well-being, engagement, and overall satisfaction (6). A study in urban planning and psychology demonstrated how physical surroundings influence human behavior, cognition, and emotional states (7). Therefore, it is reasonable to postulate that the built environment of a college campus may similarly impact students' academic engagement, and consequently, their likelihood of staying enrolled.

To investigate the influence of the built environment on college retention rates, we propose an innovative approach leveraging computer vision techniques. Specifically, we assessed the built environment using images obtained from Google Street View. Google Street View is a digital mapping

JOURNAL OF
EMERGING INVESTIGATORS

DOI: https://doi.org/10.59720/23-112

service that provides 360-degree panoramic street-level imagery, providing us with a cost-effective and comprehensive resource to access a vast collection of visual data from various college campuses. We chose to use Google Street View imagery since it met our image criteria, which consisted of good lighting, quality, angle, location, etc.

Computer vision, a field at the intersection of artificial intelligence and image processing, equips machines with the ability to extract information and gain insights from visual data. By employing image classification and computer vision algorithms, we can systematically analyze the visual features of college campuses and quantify their relationship with retention rates. Building upon the rich history of computer vision research, which has made remarkable strides in image recognition, object detection, and scene understanding, we aim to extend the application of these techniques to the domain of higher education (8-10). By doing so, we hope to contribute to a growing body of work that explores the relationship between physical environments and student outcomes.

By shedding light on the often-overlooked built environment in relation to college retention rates, our study contributes to the ongoing conversation surrounding student success in higher education. We hypothesize that one way that colleges could enhance their retention rates is by improving the aesthetic qualities of their campus surroundings. We found that there was a correlation between the Google Street View Images of the colleges and its retention rates with a 99% confidence interval. By employing computer vision techniques and leveraging the vast resources of Google Street View, we have begun to unravel the potential impact of physical environments on student outcomes. This research sets a foundation for future investigations and emphasizes the importance of considering the built environment when designing and managing college campuses.

RESULTS

Data Collection

We assembled a list of 35 colleges and universities, including both private and public institutions with various retention rates. Some of the colleges with lower retention rates include Southeastern Oklahoma State University with a rate of 0.61 and Virginia Union University with a rate of 0.68, while higher retention rate colleges included Yale and Northwestern Universityeach with retention rates of 0.99 (**Table 1**).

College	Retention Rate	College	Retention Rate
Southeastern Oklahoma State University	0.61	Santa Clara University	0.94
Virginia Union University	0.68	Emory University	0.95
University of Central Missouri	0.712	Stevens Institute of Technology	0.95
Penn State Berks	0.75	United States Military Academy at West Point	0.96
Old Dominion University	0.76	Wesleyan University	0.96
California State University, Dominguez Hills	0.78	Worcester Polytechnic Institute	0.96
Franklin Pierce University	0.79	University of Virginia	0.97
Mount Holyoke College	0.81	Carnegie Mellon University	0.97
University of Missouri–St. Louis	0.82	Vanderbilt University	0.97
John Jay College of Criminal Justice	0.83	Johns Hopkins University	0.97
Allegheny College	0.84	Swarthmore College	0.97
Munster Technological University	0.85	Stanford University	0.98
University of Cincinnati	0.86	Claremont McKenna College	0.98
University of New Hampshire	0.86	Harvard University	0.98
College of Wooster	0.87	University of Notre Dame	0.98
Creighton University	0.91	Northwestern University	0.99
University of South Florida	0.91	Yale University	0.99
Haverford	0.94		

Table 1: List of the 35 colleges with corresponding freshman retention rates. The 35 colleges used for our data and their corresponding retention rates arranged in order from least to greatest. Based on their retention rate, colleges were divided into three groups of low (0.61-0.85) shown in red, medium (0.86-0.96) shown in yellow, and high (0.97-1.00) shown in green, and given a label for machine learning prediction.

Color Detection Results

Studies have been conducted which display a positive correlation between the amount of greenery around a school and cognitive development in children (11). Because of this, our first idea was to make a model which explored if any colors were more or less present in colleges with high retention rates. We first used the K-means cluster algorithm from scikit-learn to extract a set of the four most prominent colors in each image (**Figure 1**). Computers understand colors using the RGB color model, which is used to combine different concentrations of red, blue, and green to quantify any color. Throughout our dataset of 3,003 images, we had a total of 11,456 unique RGB sets corresponding to unique colors. The most popular colors in the images from the colleges with lowest retention rates were relatively lighter colors, such as a sky blue, while the images from the medium and high colleges tended to have colors with darker hues, such as a dark green and a deep black (**Table 2**).

We then trained a decision tree model from the open-source scikit-learn python library. A decision tree model is a type of supervised machine learning algorithm used to classify and make predictions based on a labeled training set. However, the lack of measured difference in the most recurring colors between each of the three retention rate classes led to a confused model with a low accuracy of just 38.04% when a Decision Tree Classifier was built on these extracted colors (**Figure 2**). We then calculated the F1 scores of each trisection of schools to measure the precision of the model. The model did tend to perform slightly better

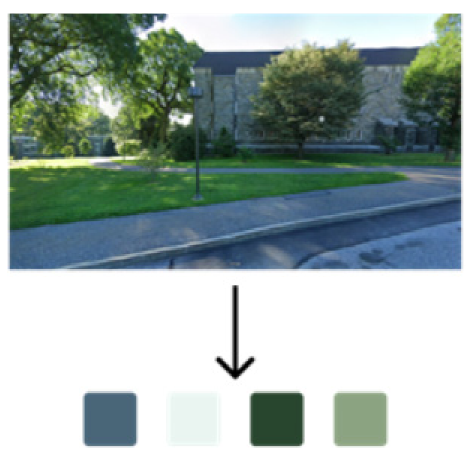

Figure 1: K-means clustering color palette extraction. Example image alongside its extracted color palette. A Google Street View image from Swarthmore College at longitude 39.9057775, and latitude -75.3530016 with camera 87° up from the ground and 38° east of south alongside the RGB colors representing the 4 clusters derived through the K-means clustering algorithm.

DOI: https://doi.org/10.59720/23-112

Retention Rate Tertile	Most frequently occurring color	2nd most frequently occurring color	3rd most frequently occurring color
Low	RGB (188, 189, 191)	RGB (51, 60, 41)	RGB (202, 208, 210)
Medium	RGB (199, 200, 201)	RGB (46, 44, 40)	RGB (42, 49, 32)
High	RGB (46, 49, 35)	RGB (1, 1, 1)	RGB (236, 239, 242)

Table 2: Most frequent colors in the images across each tertile of retention rates. Background cell color represents the color code in that cell for visualization. The three most frequently occurring colors, which were extracted from our color detection method, in each retention rate tertile. Overall, the colors, mainly light hues of gray and darker shades of green, were largely similar between each tertile.

in correctly predicting the class of retention rates when presented an image from the middle and higher classes, with a F1 score of 36.99% for the lowest retention rate class, a F1 score of 38.31% for the middle retention rate class, and a F1 score of 38.62% for the highest retention rate class.

Object Detection Results

We then wanted to see if the number of objects in these Google Street View images were correlated with the retention rates. We used the You Only Look Once version 4 (YOLOv4) pre-trained object detection network to detect objects in each of our images (**Figure 3**). The YOLOv4 is pre-trained to detect a set of various objects, some more common than others. The most common objects detected on campus were largely the same across all three retention rate tertiles but differences start to emerge when the frequency of the objects are analyzed (**Table 3**). The four most common objects that we found in our pictures were cars, people, trucks, and benches, which made up over 96% of the objects that were found. When these features were used to train a Decision Tree Classifier with the Gini function, the model had an accuracy of 40.49%, indicating a significant improvement over our previous method (**Figure 4**). We also calculated the F1 score for each retention rate class to see in which test images the model performed the best. For the lowest retention rate class, the F1 score was 41.37%, for the middle retention rate class, it was 45.25%, and for the highest retention rate class, it was 30.78%, highlighting the increased performance and unique characteristics of images from colleges in the lower 2 tertiles.

Additionally, after training the Decision Tree Classifier, we analyzed which features held the most importance. The frequency of the most common objects such as cars, people, trucks, benches, and stop signs had the largest influence on the classifier's predictions (**Table 4**). Overall, we determined that the YOLOv4 model, with an accuracy of 40.49%, was best suited for the prediction of college retention rate based on image analysis. Based on these results, we concluded that

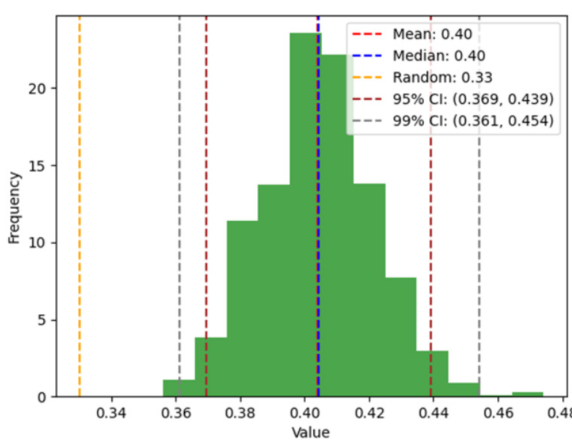

Figure 2: Accuracies of a Decision Tree Classifier trained on extracted color palettes. (n=1000). The Decision Tree Model was trained on predicting which tertile of retention rates a college would fall into based on a feature set of four colors that were extracted from sample images using a K-nearest neighbors algorithm. These accuracies were measured on a set of data portioned off for testing. The mean and median accuracies are both centered around 40%. The maroon dotted line shows the bounds of the 95% confidence interval and the gray dotted line shows the bounds of the 99% confidence interval. Complete random selection, or 33%, is shown by the yellow dotted line which is well outside of both the 95% and 99% confidence interval.

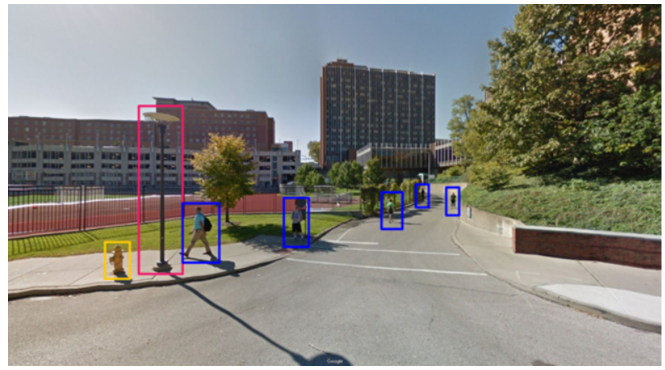

Figure 3: Detected objects in a Google Street View image from the University of Cincinnati. Boxes indicate detected objects with different colored boxes representing different types of objects. Yellow boxes represent fire hydrants, red boxes represent lamp posts, and blue boxes represent people. The YOLOv4 CNN was used to detect common everyday objects in a Google Street View image from Swarthmore College at longitude 39.1296361, and latitude - 84.5164517 with camera 90° up from the ground and 14° west of south.

Retention Rate Image Classification	Most frequently occurring object	2nd most frequently occurring object	3rd most frequently occurring object	4th most frequently occurring object
Low Retention Rate Images	car (7363)	truck (302)	person (288)	traffic light (77)
Medium Retention Rate Images	car (6127)	person (626)	truck (397)	chair (199)
High Retention Rate Images	car (5387)	person (1038)	truck (336)	bench (256)
All Images	car (18877)	person (1952)	truck (1035)	bench (488)
Percentage of Total Objects from All Images	car (81.09%)	person (8.39%)	truck (4.45%)	bench (2.10%)

Table 3: Most frequent objects in the images across each tertile of retention rates. The four most frequently occurring objects, which were extracted from our object detection method, in each retention rate tertile. Cars were the most frequently occurring object in each retention rate tertile, followed by people and trucks.

DOI: https://doi.org/10.59720/23-112

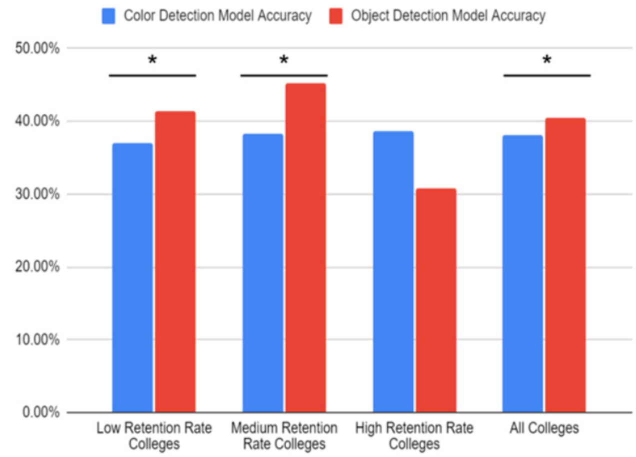

Figure 4: Accuracies of Decision Tree Model across the three tertiles of colleges trained on detected colors compared with detected objects. (n=1000). The Decision Tree Model was trained on predicting which tertile of retention rates a college would fall into based on a feature set of four colors that were extracted from sample images using a K-nearest neighbors algorithm and the objects detected in each image using a YOLOv4 CNN. * $p < 0.001$, two sample t-test.

there was a correlation between campus-built environments and the retention rates of colleges. We proceeded to further analyze our results to gain more insights into the specific objects that were associated with higher retention rates.

DISCUSSION

It is important to evaluate the correlation between campus environment and retention rate so colleges can consequently assess the relative importance of all the factors that go into making a school a place that students want to come to and stay at. Machine learning techniques can be used to draw correlations and trends between environmental factors and retention rates of colleges, making them ideal to apply in this context. Using two different methods, color detection and object detection, we reached the conclusion that there is a correlation between the external environment of a college campus and the retention rate of said college.

Machine learning relies on extensive and reliable data, so we needed a large and trustworthy dataset to work with. At first, we looked towards pre-collected datasets, like those on Kaggle or Data.world, but soon found the colleges represented lacked diversity or didn't meet the criteria that we needed. For example, we found a dataset on Data.world that met most of our needs, except for the fact that it only examined colleges in Washington DC, and it gave birds-eye-view images (12). To get the variety in data we were looking for, it became clear that we needed to make our own dataset. After realizing this, we looked at other studies that relied on images of public places, such as the Stanford study that examined images of areas in Oakland and attempted to predict the poverty rates based on the cars in the images (13). From this study, we got the idea of scraping images from Google Street View, which both allowed us to collect a random and diverse dataset, while also keeping the quality consistent.

Next, we started thinking of what makes up an environment. Two of the most important characteristics we came up with were the colors in the environment and the actual objects present in the environment. To minimize any sort of error or bias, we decided to try both approaches of characterizing the environment. Based on our results for both approaches, we reached the conclusion that there is a correlation between the built environment of a college campus and the retention rate of the school. More specifically, the object detection model showed that a high number of cars, trucks, and people were correlated with a high retention rate. This could indicate that there is a correlation between a college's outdoor environment and its retention rate.

One limitation of our findings was the small number of schools sampled. Although we did have many images from each school showing different aspects of the built environment, in total, we had a very limited number of schools in our dataset. To train the model, we used an 80/20 split, using 80 percent of data to train the model and 20 percent of data to test the model. This meant 28 out of the 35 colleges we had in our dataset were used to train the model, leaving us with only 7 different schools for testing. Even with the thousand training and testing iterations we ran, the seven schools tested lead us to see rudimentary trends in the results. As we increased the size of the training set, the mean accuracy of the model increased, which indicates that more training data would have been beneficial to the model (**Figure 5**).

Another limitation was that we only looked at two different aspects of the built environment: color and objects present. In reality, our built environment encompasses everything around us, or all the human-made conditions that surround us. Although both tested aspects are relatively broad, the fact that we don't include other aspects of the environment could contribute to the lower level of correlation found in our study. For example, we don't consider the interior of any buildings in the campus, which is important because that is where most of the learning occurs. Other studies were able to quantify the beauty of an outdoor environment by also considering whether structures were manmade or natural or the amount of sunlight, which is a way to proceed with this experiment in the future (14). Overall, further research into the importance

Rank	Object	Importance
0	car	0.382939
1	truck	0.119193
2	person	0.112441
3	bench	0.068939
4	stop sign	0.056521
5	bicycle	0.029856
6	bus	0.029624
7	potted plant	0.028526
8	chair	0.027963
9	traffic light	0.022505
10	train	0.017148

Table 4: Features with most importance in Decision Tree Classifiers trained on detected objects. The importance of various objects extracted from our object detection method. The importance shows the significance of each object during the Decision Tree Classifier's prediction process. Cars, followed by trucks and people, were the most influential for the Decision Tree Classifier when making a prediction.

JOURNAL OF
EMERGING INVESTIGATORS

DOI: https://doi.org/10.59720/23-112

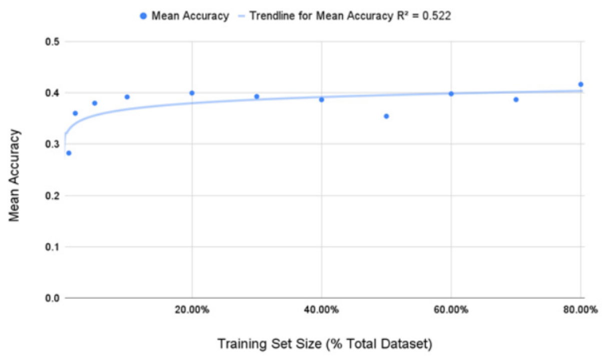

Figure 5: Accuracies of Decision Tree Models trained with varying sizes of data. Mean accuracies (n=1000) and a logarithmic regression trend line. Decision Tree Models were trained on predicting the tertile of retention rates a college would fall into based on a feature set of the frequency of certain objects detected in each image using a YOLOv4 CNN. There is a strong logarithmic correlation (R^2 = 0.552) between the amount of data used to train a model and the accuracy of the resulting trained model.

of different parts of the environment could help solve this problem.

Furthermore, we were limited by the quality of the images that we selected. An image that was taken with low quality could impact both models that we tried. For example, a bad image would make it much harder for the models to detect certain objects or it would make colors muddier, completely changing the shade, affecting our color detection model. Furthermore, finding a way to increase the resolution of an image could allow the model to detect objects better and therefore increase our accuracy with the object detection model.

Various studies were able to use Google Street View more efficiently. For example, a study using Google Street View to predict demographic makeups used the same YOLOv4 model that we used to identify only cars (15). From there, they used other algorithms to determine the model and make of these cars, their price, and their fuel range, and used publicly available information on political stances, race makeup, and percentage of people with high school degrees in each city, which offered a more complete approach to their research (15). Another study was successful in predicting student retention rates using other features, such as previous academic performance, socio-economic factors, or relationships in their family life, which provided a more holistic approach to understanding the factors that come into play when predicting retention rates that our model simply wasn't testing (16).

Overall, our findings, even with their limitations, show that there is a correlation between what students are physically surrounded by and the rates at which they stay at school. This could be used by schools to help set their priorities in terms of what they want to focus on to improve retention rates. Since our research only shows a correlation and nothing more, it would take further studies to figure out what exactly colleges should do and how much of it they should do. Our research, rather than being the sole tool for colleges to use while making decisions, should be a supplementary tool for colleges that shows the outside environment should be one of the many considered features while deciding where to spend their time and money to keep students learning.

MATERIALS AND METHODS

Data Collection

A list of 35 colleges and universities was assembled, including both private and public institutions with various retention rates. Images of the outdoor environments of these colleges were sourced from Google Street View. For each of the 35 colleges, a list of 80-100 different latitude-longitude points was created, ensuring that they were spread out across the campus and covered different areas of interest. A camera angle for each coordinate was randomly selected and a standardized 2560x1440 image was downloaded using the JavaScript puppeteer library.

Retention rate data was obtained from College Factual, an organization that uses various sources to create an estimate of the true data of a college. Colleges were split into three evenly-sized tertiles– one for retention rates between 61% and 85%, one for retention rates between 86% and 96%, and another for retention rates above 96%. Each image was assigned a label representing the tertile of the college that the image was taken at.

Color Detection Method

After images were collected from Google Street View, the goal was to extract information about the built environment through machine learning image recognition techniques. Our initial approach was to assess the ambiance in colors. To do this, the RGB values were extracted for each pixel in the image, resulting in a matrix of red, green, and blue values for each image. K-means clustering was then used to create a color palette of four colors for each image. The clustering was applied to the matrix of RGB values and assigned each pixel to one of four clusters, each represented by its centroid color. Finally, a decision tree classifier was trained on the four-color clusters for each image.

Object Detection Method

To improve the accuracy of our analysis, we moved away from color detection and focused on detecting the actual objects in the images. Understanding that training our own CNN to detect objects present in the images was difficult and resource-intensive, exploring the use of pre-trained object detection models that were already available, specifically YOLOv4, seemed suitable.

The YOLOv4 model is a state-of-the-art deep learning model that was developed in 2020 (17). It is an improvement over previous versions of the YOLOv3 model and is considered one of the most accurate and efficient object detection models available, giving a groundbreaking result of 65.7% AP50 accuracy on the MicroSoft COCO dataset (17). This model uses a single neural network to predict the bounding boxes and class probabilities for objects in an input image. It does this by dividing the input image into a grid and predicting the class and location of each object within each grid cell. The model then outputs a list of bounding boxes and class probabilities for each object detected in the image. The YOLOv4 model is trained on a large dataset and is capable of detecting over 80 different object categories with high accuracy. It is also designed to be highly efficient, allowing it to process images in real-time on a standard GPU.

It was decided to use the YOLOv4 model for our project because of its high accuracy and efficiency in object detection as we had a large set of images to process. Using

DOI: https://doi.org/10.59720/23-112

a pre-trained model like YOLOv4, time and resources were saved that would have been required to train our own object detection model from scratch. With this pre-trained model, a script in python was developed to detect objects in each image. A decision tree classifier was then trained to predict the retention rates of each college based on the detected objects in the corresponding images by utilizing the object detection results.

Statistical Analyses

To perform statistical analyses, a decision tree model was selected to find correlations between two sets of data by passing the input data, known as features, through a set of nodes to eventually make a decision and predict a label for that input.

Individually, for both our data extracted from the images using the color detection method and from the object detection method, a decision tree model was trained using scikit-learn's DecisionTreeClassifier (version 1.3). Gini Impurity was used to measure the quality of a split and choose the "best" split at each node, with an unlimited maximum tree depth, a minimum number of 2 samples required to split an internal node, and a minimum of 1 sample required to be at a leaf node. After splitting the dataset into 80% for training and 20% for testing, the model could be continued to be trained.

Using the data extracted by the color detection method described above, a decision tree model was trained, and the resulting accuracy and F1 scores were recorded. This was repeated 1000 times. The mean and standard deviation of the dataset were calculated in Python scikit-learn to determine if there is a significant correlation between the colors given and the retention rates. Using the mean and standard deviation, was not in the 99% confidence interval, it can be concluded with 99% certainty that there is a correlation between the outdoor environment of the colleges and its retention rate.

Received: May 23, 2023
Accepted: August 3, 2023
Published: January 02, 2024

REFERENCES

1. National Center for Education Statistics. (2022). Undergraduate Retention and Graduation Rates. *Condition of Education*. U.S. Department of Education, Institute of Education Sciences. nces.ed.gov/programs/coe/indicator/ctr. Accessed 9 July. 2023.
2. "Graduation and Retention Rates - What is the full-time retention rate in postsecondary institutions?" National Center for Education Statistics. nces.ed.gov/ipeds/TrendGenerator/app/answer/7/32. Accessed 14 May. 2023.
3. Lotkowski, Veronica A., et al. "The Role of Academic and Non-academic Factors in Improving College Retention: ACT Policy Report." PsycEXTRA Dataset, 1 Jan. 2004, https://doi.org/10.1037/e420492008-001.
4. DeNicco, James, et al. "Factors of one-year college retention in a public state college system." Research in Higher Education Journal, vol. 27, Jan 2015, pp. 3-5
5. Gallagher, Strawberry. "The Impact of Resources and Programs on EOF Students: A Study to Improve On-Campus Support Programs." 2023. Saint Elizabeth University, PhD dissertation.
6. Wen, Qiang et al. "Evaluation of Satisfaction with the Built Environment of University Buildings under the Epidemic and Its Impact on Student Anxiety." International journal of environmental research and public health vol. 20,5 4183. 26 Feb. 2023, https://doi.org/10.3390/ijerph20054183.
7. Hartig, Terry, et al. "The Need for Psychological Restoration as a Determinant of Environmental Preferences." Journal of Environmental Psychology, vol. 26, no. 3, Sept. 2006, pp. 215–26, https://doi.org/10.1016/j.jenvp.2006.07.007.
8. Xin Feng et al. "Computer vision algorithms and hardware implementations: A survey."Integration, vol. 69, 2019, p. 309-320, https://doi.org/10.1016/j.vlsi.2019.07.005.
9. Bhil, K. et al. (2022). "Recent Progress in Object Detection in Satellite Imagery: A Review." In: Aurelia, S., Hiremath, S.S., Subramanian, K., Biswas, S.K. (eds) Sustainable Advanced Computing. Lecture Notes in Electrical Engineering, vol 840. https://doi.org/10.1007/978-981-16-9012-9_18.
10. Ji Hou et al. "Exploring Data-Efficient 3D Scene Understanding With Contrastive Scene Contexts" Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), 2021, pp. 15587-1559
11. Dadvand, Payam et al. "Green spaces and cognitive development in primary schoolchildren." Proceedings of the National Academy of Sciences of the United States of America vol. 112,26 (2015): 7937-42. https://doi.org/10.1073/pnas.1503402112.
12. DC, Code For. data.world/codefordc/college-university-campuses. Washington DC, USA, 4 Oct. 2017.
13. Acharya, Ambika. "Neighborhood Watch: Using CNNs to Predict Income Brackets from Google Street View Images." Stanford Vision, 2017.
14. Seresinhe et al. "Using deep learning to quantify the beauty of outdoor places." Royal Society open science 4.7 (2017): 170170. https://doi.org/10.1098/rsos.170170
15. Gebru, Timnit, et al. "Using deep learning and Google Street View to estimate the demographic makeup of neighborhoods across the United States." Proceedings of the National Academy of Sciences 114.50 (2017): 13108-13113. https://doi.org/10.1073/pnas.170003511
16. Matz, Sandra C., et al. "Using machine learning to predict student retention from socio-demographic characteristics and app-based engagement metrics." Scientific Reports 13.1 (2023): 5705. https://doi.org/10.1038/s41598-023-32484-w
17. Bochkovskiy, Alexey, Chien-Yao Wang, and Hong-Yuan Mark Liao. "Yolov4: Optimal speed and accuracy of object detection." arXiv preprint arXiv:2004.10934 (2020). https://doi.org/10.48550/arXiv.2004.10934