



Seeing Floods from Space: How Sentinel-1 Helps Map Floods in New York City

Rishikesh Dasgupta^{1,2,*}; Tashfia Dih^{1,3}; Clare Hill^{1,4}; Saroni Sikder^{1,5}; Yuneeb Awan⁶; Ali Haider⁷ and Reza Khanbilvardi⁷

Submitted: 12 October 2025 Accepted: 15 October 2025 Publication date: 29 October 2025

DOI: [10.70671/bce5na79](https://doi.org/10.70671/bce5na79)

Abstract: Flood susceptibility maps are highly effective tools for identifying and mitigating areas at risk of flooding, but they require extensive historical flood data from either in situ or remote sensing data collection. This study demonstrates how publicly available Sentinel-1 Synthetic Aperture Radar imagery can be used to generate flood maps for New York City's five boroughs. Sentinel-1 provides high spatial, geometric, and radiometric resolution, making it ideal for detecting surface water under all weather conditions. Using the European Space Agency's Sentinel Application Platform for data preprocessing and QGIS and Google Earth Pro for visualization, we processed and analyzed Sentinel-1 data to identify flood-prone zones. The results highlight Sentinel-1's effectiveness as a low-cost, reliable resource for urban flood analysis and demonstrate its potential for supporting flood monitoring and resilience planning in major metropolitan areas like New York City.

Author keywords: Urban flooding; Sentinel-1; SAR; Remote sensing; New York City

Introduction

Urban flood risks in New York City

Urban flooding is the result of the inflow of stormwater surpassing the capacity of a drainage system to permeate water into soil or carry it away.¹

Arguably, the main cause of urban flooding is precipitation, specifically heavy rainfall, which has the greatest impact on urban flooding.² Rainfall intensity is vital in determining the design flow rate, a key parameter when designing drainage structures. Rainfall intensity is often accompanied by formulated relationships drawn between rainfall intensity, storm duration, and return period in intensity-duration-frequency curves. NYC's exposure to pluvial and coastal flooding is illustrated in Fig. 1.

Another cause of urban flooding is topography, particularly slope and elevation.² There tends to be an increase in water velocity and discharge with larger slope angles. Consequently, there is a lower mean water depth and peaks in stored runoff in areas with larger slope angles, which explains why flatter surfaces are at a higher risk of flooding. Areas low in elevation are more vulnerable to flooding. For example, urban areas 10 meters above sea level have a flood risk of 1.3%, whereas urban areas 10 meters below sea level have a flood risk 3.77 times higher than that. Areas close in elevation to the ocean or a lake are also vulnerable to storm surges, sea level rise, or regional precipitation, which make them more prone to flooding. Notably, the rate of global mean sea level rise has increased from ~2.1 mm/year in 1993 to ~4.5 mm/year in 2023,³ which increases the potential for urban areas near the coast to flood.

Land features and infrastructure also have an impact on urban flooding.² Engineered impervious surfaces, such as sidewalks or buildings, increase the risk of flooding. Imperviousness and drainage constraints are summarized in Fig. 2. The runoff coefficient is the infiltration ability of an area and it is directly related to the peak runoff rate (water discharge). Impervious surfaces have a runoff coefficient between 0.70 and 0.95, whereas green surfaces like grass and soil range between 0.05 and 0.35. The higher runoff coefficient of impervious surfaces and the direct relationship between the coefficient and rate suggest that there is higher discharge for impervious surfaces. In particular, buildings prevent infiltration, increase runoff, and affect where and how the water flows. As a result, catch basins are crucial for rainwater to enter the sewer system. When clogged, they can be a primary

*Corresponding Author: Rishikesh Dasgupta. Email: rishikeshdasgupta800@gmail.com

¹CUNY CREST High School Initiative in Remote Sensing of Earth System Engineering and Sciences (HIRES), The City College of New York 10031

²John Adams High School, New York 11417

³Stuyvesant High School, New York 10282

⁴The High School of Art and Design, New York 10022

⁵The Bronx High School of Science, New York 10468

⁶Earth and Atmospheric Sciences, The City College of New York 10031

⁷Department of Civil Engineering, CUNY-CREST Institute, and United Nations University (UNU) Hub on Remote-Sensing and Sustainable Innovations for Resilient Urban Systems (R-SIRUS)-UNU Institute for Water, Environment and Health (UNU-INWEH), The City College of New York 10031

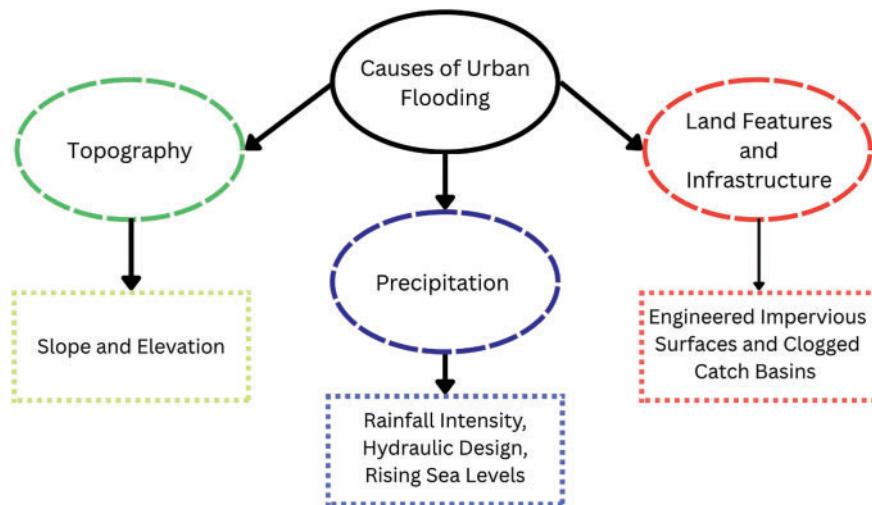


Figure 1. Causes of urban flooding

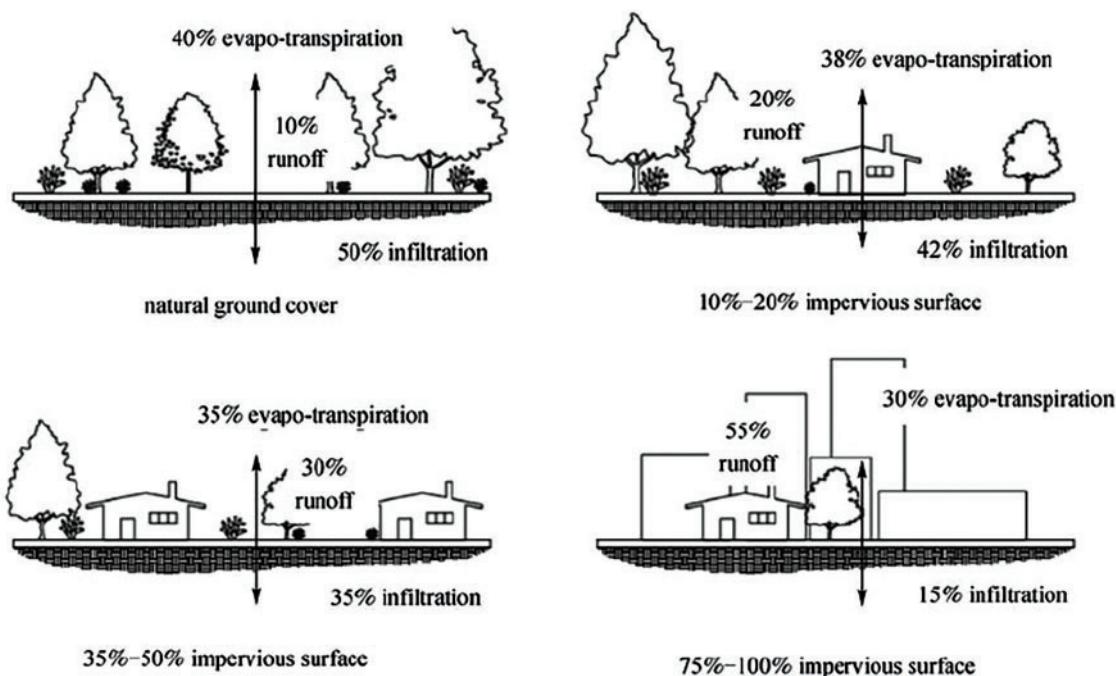


Figure 2. Impervious surface changes due to urbanization and its impact on local hydrology²

indicator of an incoming flash flood, making areas with less-maintained catch basins more susceptible to flash floods. Clogged catch basins remain a prevalent problem in New York City (NYC), with significant complaints reported in nearly half (47%) of the city's ZIP codes between 2010 and 2019.²

Need for research and impacts of urban flooding

In New York, the sea levels have risen 304.8 mm since 1900 and are projected to increase by as much as 1645.92 mm by 2100, leading to increased flooding. By the 2050s, a storm like Hurricane Sandy, with wind gusts of up to 80 mph at landfall and surging rainfall that the city was not prepared for, could cause \$90 billion in economic damage, which

would be nearly 5 times Sandy's impact.⁴ Sea level rise for New York from 2020 to 2050 is shown in Fig. 3, derived from the U.S. Sea Level Change (2024)⁵ projections.

Two particularly vulnerable groups to urban flooding are elderly people and low-income families.^{6,7} Elderly people and those with limited mobility are the most vulnerable to coastal storm-related injuries and fatalities, considering how crucial a swift evacuation becomes during a flash flood.⁶ Repairing a home or replacing belongings can place a significant financial strain, especially on low-income families. For instance, the damage from Hurricane Ida, during which 3.75 inches of rain fell per hour, was concentrated in Queens, Brooklyn, and the Bronx, which are populated with lower-income and immigrant communities.⁷ More than a dozen people died because of basement flooding.

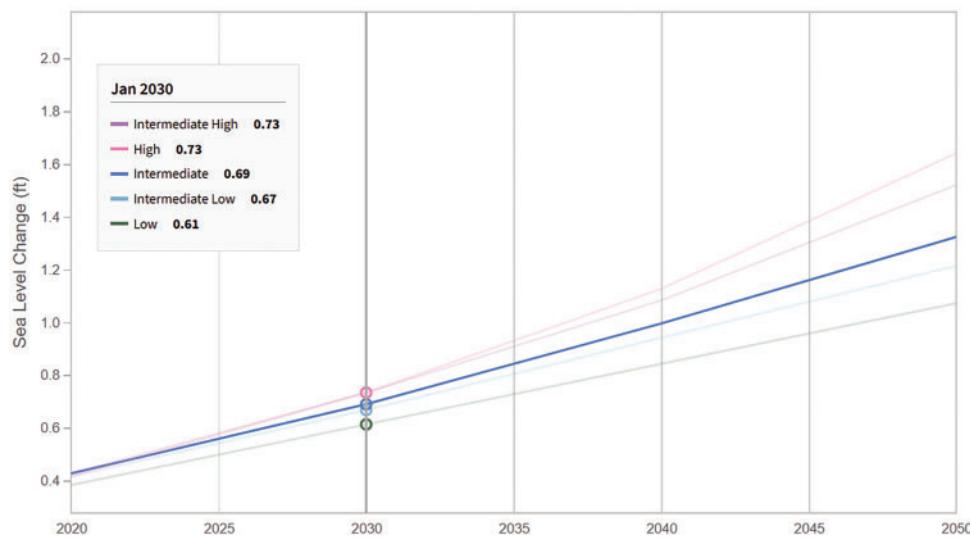


Figure 3. Sea level rise for New York from 2020 to 2050⁵

There are two methods of collecting rainfall data: in situ and remote sensing. In situ involves a direct measurement instrument, namely, a rain gauge.² The instrument error of a rain gauge decreases with rainfall intensity, but the distance between rain gauges is problematic, particularly for small-scale urban drainages, as the spacing issue can contribute to up to 20% of total uncertainty. Although remote sensing has certain disadvantages, such as cloud top reflectance and thermal radiance, using radar and satellite data can provide sufficient spatial distribution for rainfall estimations.

The substantial impacts of flooding and the huge risk of flooding in NYC present a need to find and create tools to prevent and map flooding in NYC. Flood susceptibility maps (FSMs) show the probability of flooding in a particular area and are one of the most effective strategies for flood prevention and mitigation.³ Identifying the appropriate method, resolution, and parameters of data sources is crucial to obtaining reliable results. A crucial component is the analysis of past flood data to identify which areas are vulnerable to floods and to pinpoint potential improvements that can be made in areas with similar characteristics, such as slope angles and elevation.

Although physical mapping is essential in urban flood research, there is a contemporary incorporation of statistical and machine learning techniques to analyze patterns in historical data and make predictions.² Instead of competing these methods against each other, it is beneficial to utilize them together to solve intricate environmental problems.⁴

Potential solutions

Engineers and employees in public policy alike have proposed improvements to the city's infrastructure that would minimize the risk of severe flash flooding in vulnerable regions, thereby reducing detrimental costs from damages, and protecting the public health of NYC residents.

The combined sewer systems currently implemented in New York are able to hold 1.75 inches of rainwater and sewage water per hour, where both rainwater and sewage water are transported to a sewage treatment plant via a single pipe, see Fig. 4 for a general illustration of a combined sewer and runoff management.⁷ However, when faced with heavy rainfall during severe weather events, the sewers are unable to handle the excessive water. As a result, this combined system makes flash floods and the contamination of water more likely.⁹ Based on this, the Extreme Weather Response Task Force of NYC writes that it is vital to invest in an upgraded sewer system that has separate pipes for rainwater and sewage, also known as a municipal separate storm sewer system (MS4), particularly in high-risk areas with low elevation.¹⁰ Although this upgrade process can cost up to 100 billion dollars to effectively manage and complete, these infrastructural changes can drastically improve NYC's defense against severe weather, making advocacy a crucial factor in investing in sewer infrastructure.⁸

As mentioned, the materials used in NYC's infrastructure development tend to be impervious surfaces that make it difficult for water to absorb into the surface without manifesting in a flash flood. Plus, impervious surfaces encourage the urban heat island effect because they absorb more heat, which results in hotter surface temperatures.¹¹ Combating the issue of the urban heat island effect in NYC and other urban areas alike remains a crucial aspect of protecting the city from flash flood events. The lack of vegetation in parts of NYC is also a major contributor to the urban heat island effect. A way to combat both flash flooding and the urban heat island effect has been the implementation of rain gardens that would work to collect and manage stormwater that are on the sidewalks.¹² However, these rain gardens are currently limited in development and may not cover all areas that are susceptible to flooding, especially when looking at locations near low-elevation shorelines.¹² This connects to socioeconomic segregation because, as evidenced by multiple

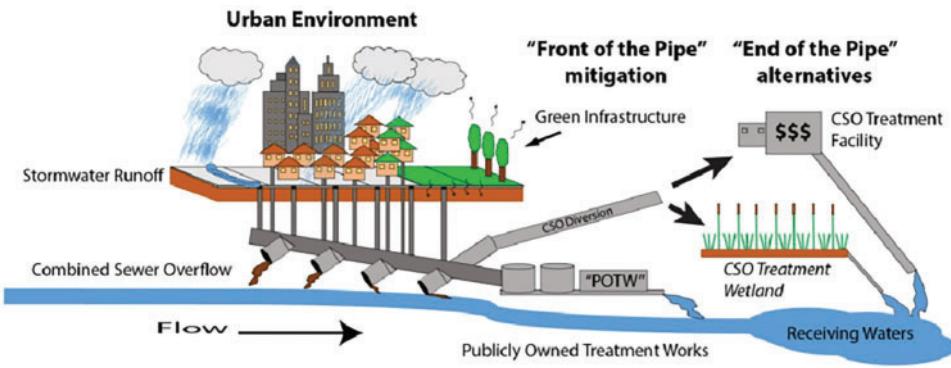


Figure 4. Process diagram illustrating urban runoff and combined sewer overflow management¹³

urban areas, lower-income residents are more susceptible to the urban heat island effect alongside flash flooding.¹⁴

Government officials have made strides to invest in flood warning systems across NYC that would prevent fatalities and injuries and protect assets from being damaged during a flash flood.¹⁵ In contrast, a flood with a delayed warning would cause more damage to people and their property. Simultaneously, providing grants and funds for flash flood research can improve the predictability of flash floods caused by severe weather.¹⁶ As previously mentioned, FSMs can be used to directly identify high-risk flood areas based on previous outcomes of severe weather. With this data, infrastructure improvements that prevent flooding can be concentrated in those areas, benefiting the city both economically and humanely.

Description of Sentinel-1

Sentinel-1 (Fig. 5), according to the GMES Sentinel-1 mission, is a polar-orbiting, two-satellite constellation that routinely operates all day via conflict-free, 6-day exact operations with a medium resolution of 10 m for observations.¹⁷ The majority of its products aim to improve the quality of emergency response systems, marine surveillance, sea-ice monitoring, and the detection of subsidence and landslides. The products of Sentinel-1 are publicly available through the Copernicus Browser in Level-0 (raw data), Level-1 (complex SLC imagery for interferometric applications + ground range detected geo-referenced imagery, GRD), and Level-2 (geolocated geophysical products, specifically ocean products with data including wind and waves), within four modes of operation and four polarization channels (Sentinel-1 Burst ID Map).

As shown in Fig. 6, the four different modes of operation found within Sentinel-1 data collection vary in swath width and geometric resolution. The most common is Interferometric Wide-swath mode (IW) with a 250 km swath width and a high geometric resolution of 5 m × 20 m. It is equipped with a ScanSAR mode for progressive azimuth scanning and TOPSAR (Terrain Observation by Progressive Scan) to harmonize performance and reduce the radar scalloping phenomenon.¹⁸

Wave mode (WV) captures vignettes of 20 km × 20 km size with a 5 m × 5 m ground resolution in intervals

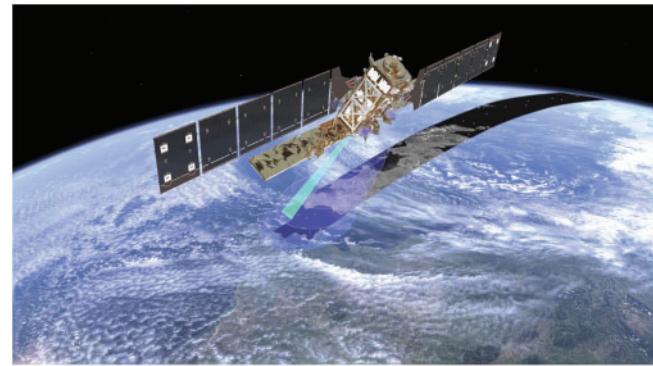


Figure 5. Sentinel-1 detects maritime traffic¹⁹

of 100 km, with a low bit rate through a unique single-polarization mode (HH or VV), while others use dual polarization. Strip Map mode maintains a 5 m × 5 m resolution and a swath of 80 km, which ultimately cover access ranges of up to 375 km. During each swath, the antenna is configured to generate a beam of fixed azimuth and elevation pointing. Extra Wide-swath mode has a swath of 400 km at a medium resolution of 20 m × 40 m, which is, once again, equipped with ScanSAR fast beam elevation scanning and TOPSAR. For the purposes of flood research in NYC, using Sentinel-1 Level-1 GRD data, IW operational mode, and VV + VH amplitude polarization channels will be used.

Sentinel-1 works by measuring and reflecting Synthetic Aperture Radar (SAR) energy to determine the radar cross section in decibels, which can be used to identify land cover types and severe weather events like floods in a process called interferometry, where waves are overlapped and information is extracted. The satellite requires internal and external calibration, culminating in more accurate satellite imagery, making it an extremely high-resolution, high geometric, and high radiometric resolution satellite. Internal calibration follows a process that sends a radar pulse to its own receiver and assesses changes in amplitude in order to readjust itself. External calibration defects outliers and dictates proper radar signaling for those areas. If there is an inconsistency in radar signaling, its constant is recalculated to accommodate. With a life cycle of 15–20 years, the satellite's quality of

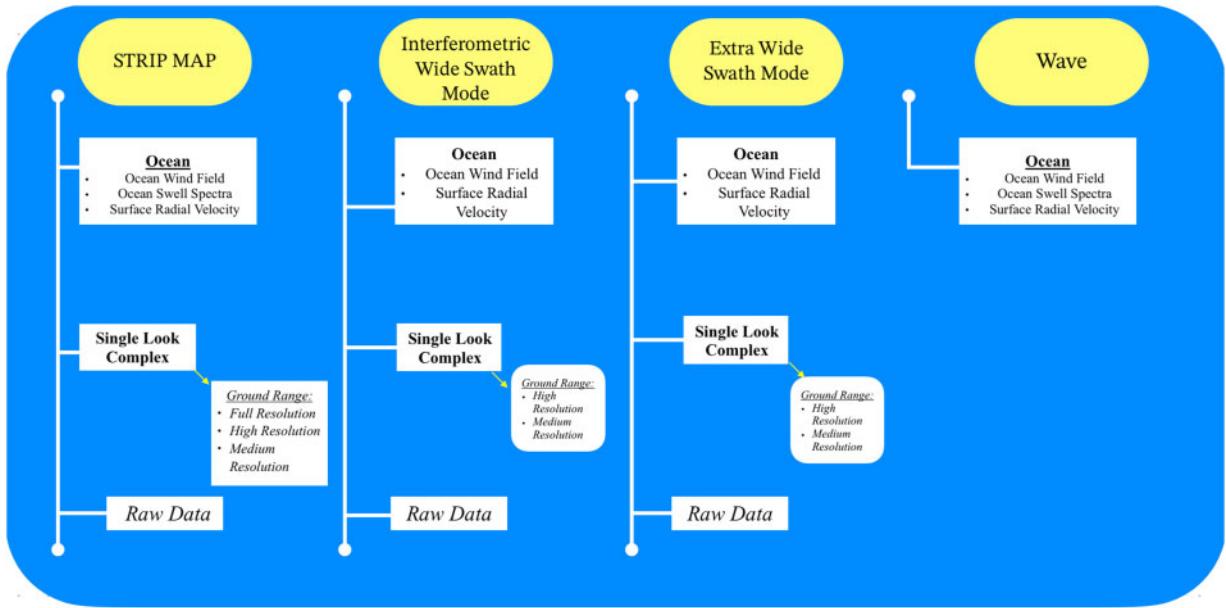


Figure 6. Sentinel-1 data products¹⁹

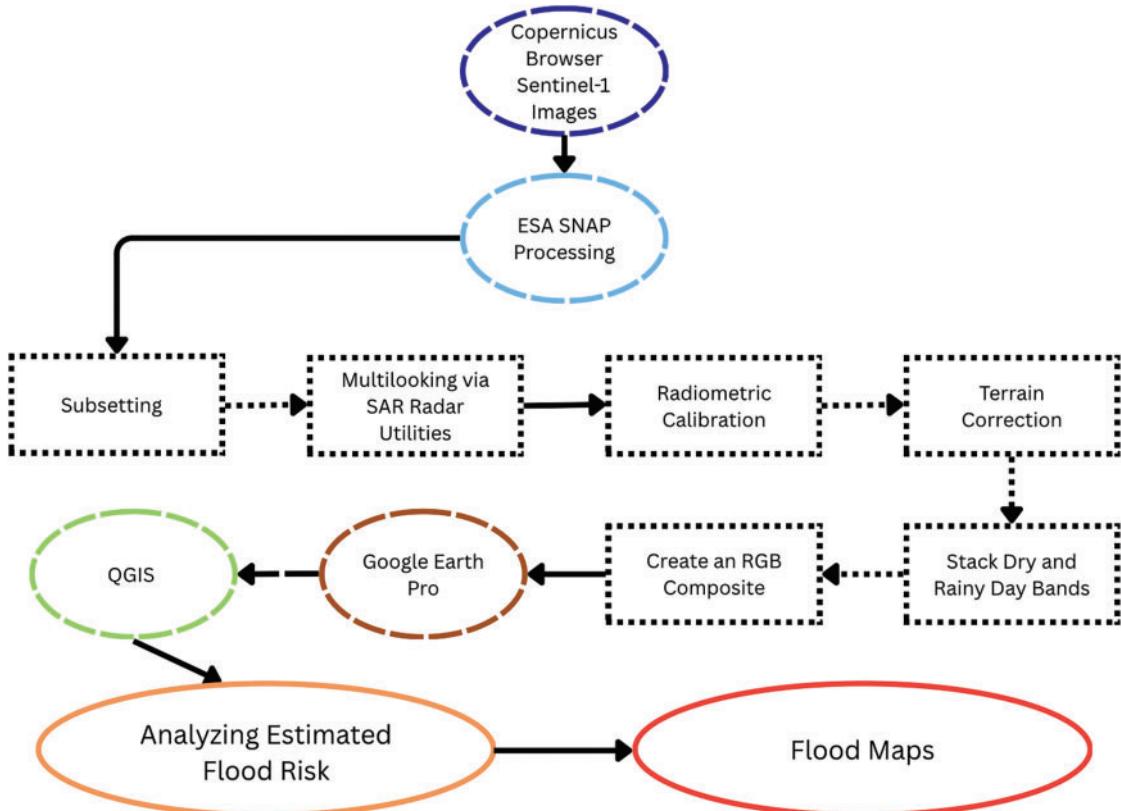


Figure 7. Methodology to create processed Sentinel-1 images

products may decay over time due to the weakening potential to send/receive signals and carry out proper calibration protocols.

For our data collection purposes, accessing the publicly available data provided by the European Space Agency (ESA) is vital to our methodology and flood research, making this aspect of Sentinel-1 unique compared to privately

owned satellites that may restrict data behind a paywall or require advanced permissions.

It is well established that Sentinel-1 data products have effectively helped researchers create flood maps for certain regions, such as Morocco.²⁰ However, these studies have not been conducted in NYC, where urban flooding is a crucial problem. The primary goal of our study was to determine

whether the Sentinel-1 satellite's data products can be used to detect urban flooding in NYC.

Methodology

Copernicus browser Sentinel-1 data

The end-to-end preprocessing workflow is presented in Fig. 7. Data collection was a matter of accessing the ESA Copernicus Browser, where the public can freely access Sentinel satellite data. For the purposes of flood research, we used Sentinel-1 data sources for its radar imagery capability, specifically Level-1 GRD. Products were to be delivered immediately upon request without further processing times for data access. Our polygonal area of interest was limited to the five boroughs of NYC, but the satellite would focus on a larger area, which would provide a satellite reading area of 43613.20 km². Time observation dates were within 1 day after the flood for the data source from March 24, 2024, where there was a total of 2.9 inches of rainfall recorded on the previous day.²¹ We picked a data source on the day of the flash flood for July 16, 2023, where there was 1.7 inches of rain, making for a significant comparison in the resulting flood outlook made visible by the Sentinel-1 satellite.²² A summary of all Sentinel-1 scenes is provided in Table 1.

Table 1. Sentinel-1 data summary

S-1 scene	UTC	EDT	Rainfall within the last 24 h
2024-03-24	22:51:30	18:51:30	2.9 in
2023-07-16	22:51:31	18:51:31	1.7 in

ESA SNAP

Sentinel Application Platform (SNAP) (version 11) was developed by the ESA²³ for visualizing Sentinel data and was a crucial tool in visualizing the floods.

Subsetting the image

We narrowed our area of interest to NYC by subsetting the initial raster GeoTIFF layer to make our data clearer and increase the speed of preprocessing.

Multilooking through SAR radar utilities

Multilooking a SAR image is crucial for reducing the speckle noise present in the image to improve the resulting visualization and analysis.²⁴ SNAP is able to carry out the multilook method through space-domain averaging, where the SAR image is a single-look complex image averaged with a small sliding window. Furthermore, selecting the number of range and azimuth looks is based on a GR square pixel, where the range look value for our purposes is 3, while azimuth looks are set to the default value of 5.

Radiometric calibration

In order for these SAR images to be acceptable for quantitative use, calibration of the image, where pixel values are related to radar backscatter, is crucial.²⁴ SNAP is able to perform this through simple radar calibration to ensure that radiometric bias is kept to a minimum, improving the visual interpretability of the SAR images. Once completed, convert the bands from linear to decibels.

Terrain correction

We performed terrain correction to correct geometric distortions due to the perspective of the sensor view angle and ground terrain.^{24,18} The algorithm traverses each grid cell in the Digital Elevation Model (DEM) to simulate the SAR imagery with the DEM. The simulated and original SAR imagery are co-registered, and a WARP function maps the simulated SAR imagery to the original SAR imagery in its corresponding position. The algorithm computes its corresponding pixel position in the simulated SAR imagery for each cell in the DEM grid. Using the WARP function, the corresponding pixel position in the original SAR image can be found, which allows the algorithm to obtain the pixel value for the orthorectified image using interpolation.²⁵

Stack the dry day and rainy day bands

To effectively see the flooded areas in our SAR images of NYC, we need to co-register the two products by stacking them. Here, we used the Create Stack operator where the flood day SAR image is resampled to share the same geographical information and dimensions as the dry day SAR image.²⁴ Immediately after, ensure that the bands are still measured in decibels.

Create an RGB composite

The final step for SNAP preprocessing is to overlay a color map that integrates three separate bands, which indicate different prominent geographic features, into a single GeoTIFF file. Band assignments and color interpretation are listed in Table 2.

QGIS

QGIS (version 3.44) is a widely used free and open-source software for visualizing geographic data. For our purposes, QGIS was used to present flood maps processed in SNAP in a concise and informative manner, with a proper title, appropriate legend, and scale, along with a compass indicating the north direction following the corrected terrain.

Results

We researched two floods, occurring on July 16, 2023, and March 24, 2024 with a substantial amount of rain, flood maps of both days are shown in Figs. 8 and 9, respectively. Table 2 outlines how each color band in the flood maps corresponds to specific surface or flooding conditions.

Fig. 10 shows a flooded portion of Staten Island, specifically Freshkills Park in Staten Island, which was converted from a landfill into a park.^{26,27} The red band represents

flooded regions 5 hours after it rained. Around \$2 million was invested in this area to convert it into a park, but Freshkills Park gets flooded easily. Fig. 10 shows a substantial part of the park being flooded along the waterways. This persistent flooding raises concerns about the long-term viability and cost-effectiveness of the investment. Each

flooding event not only damages park infrastructure but also disrupts recreational use and poses environmental risks, such as soil erosion and water contamination.

The central areas of NYC are situated at a higher elevation compared to the surrounding coastal regions. Additionally, the city has mostly impervious surfaces such as concrete,

Table 2. Band representation

Band color	Flood difference
Red band	Red indicates areas with a decrease in backscatter on the flooded day, consistent with <i>smooth water surfaces</i> detected by SAR. This decrease typically represents surface water flooding .
Blue band	Blue indicates areas with an increase in backscatter on the flooded day, consistent with <i>rough surfaces</i> detected by SAR. This increase typically represents land cover/structural changes .

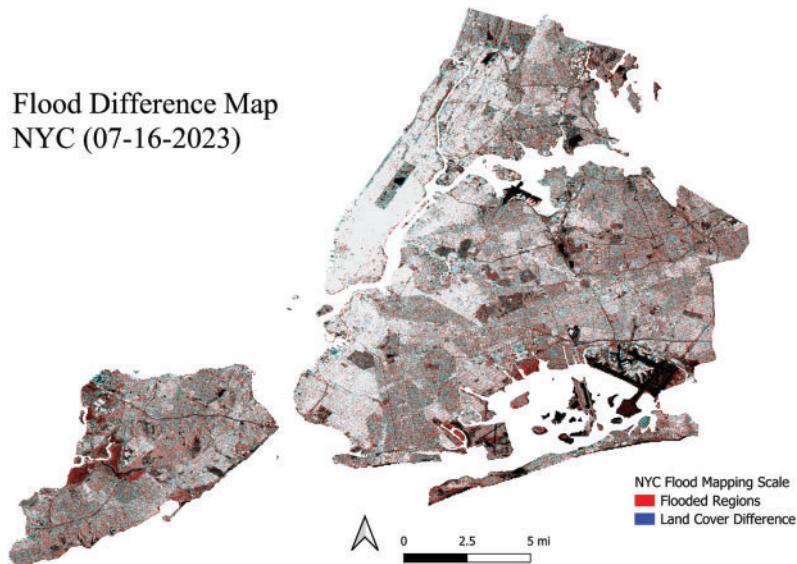


Figure 8. Flooded areas in 5 boroughs of New York City on July 16, 2023

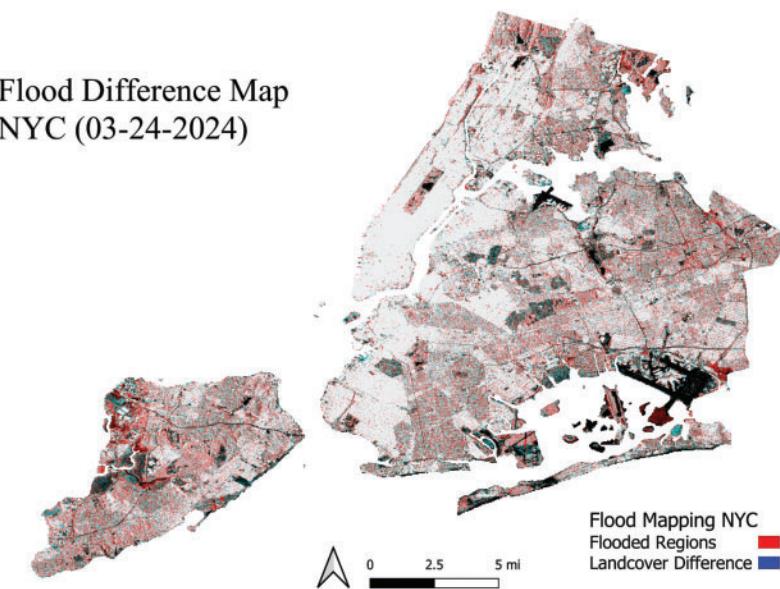


Figure 9. New York City on March 24, 2024

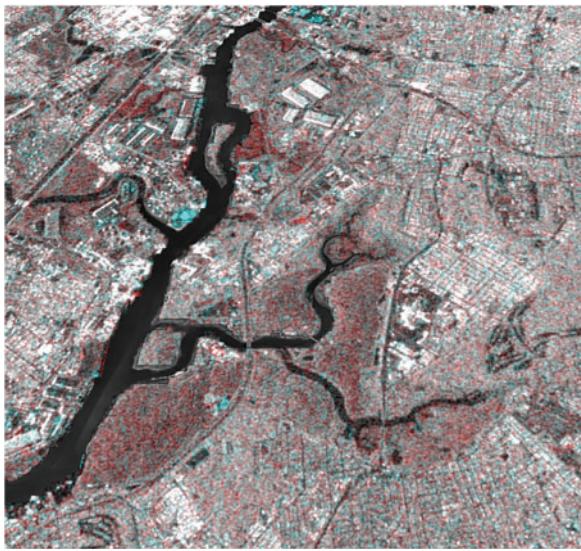


Figure 10. Staten Island on July 16, 2023

asphalt, and buildings, which result in a higher water runoff rate toward lower elevations. As previously mentioned, this runoff significantly increases the risk of flooding in coastal and low-elevation zones. About 30% of NYC is covered by an MS4 drainage system, which makes flood management in these areas more difficult. Additionally, flooding in these areas may result in water contamination and raise health concerns, as the sewage and drainage systems are combined.⁹

Fig. 9 shows flood detection results for NYC on March 24, 2024, based on Sentinel-1 imagery acquired at 6:51 PM EST on March 23, 2024. Rainfall data from March 23, 2024, indicate consistent precipitation from 5:00 AM to 3:00 PM, increasing from 0.163 inches per hour to 0.559 inches per hour and totaling approximately 2.943 inches over a 10-hour period.^{28,29} This sustained rainfall was expected to cause localized flooding; however, the SAR acquisition occurred approximately 30 hours after peak rainfall, potentially reducing visible inundation in the radar imagery as water either drained or evaporated.

Our flood detection map for March 24, 2024, reveals inundation primarily in coastal and low-lying regions such as Freshkills Park in Staten Island (shown in Fig. 11) and the Jamaica Bay area in Brooklyn. These regions exhibited pronounced decreases in backscatter intensity, suggesting the presence of residual surface water. In contrast, higher-elevation and densely built areas, including Manhattan and the Bronx, displayed minimal flood signatures, likely due to efficient drainage and rapid surface runoff from impervious materials such as concrete and asphalt, which prevented prolonged surface water accumulation detectable by SAR.

The flood maps demonstrate Sentinel-1's sensitivity to subtle surface changes in complex urban environments. These observations validate the satellite's capability to detect flooding patterns in NYC, even with delayed acquisition times. The resulting SAR flood layers generated through our processing form the basis for the subsequent comparison and validation in the Discussion section.

Discussion

We validated these satellite images signaling a flood through news articles covering the floods that occurred on that day. On March 24, 2024, it was found to be the third wettest March day in NYC history.³⁰ Additionally, roads were severely flooded, leaving drivers stranded and in need of rescue.³¹ On July 16, 2023, the rainy outlook was predicted on the same day and addressed by Kathy Hochul, with copious amounts of rain expected to fall on the city, resulting in major flooding in the borough of Brooklyn.³²

To further validate our SAR-derived flood observations, we compared the March 24, 2024, Sentinel-1 flood extent with FEMA's Preliminary Flood Insurance Rate Maps for NYC, as shown in Fig. 12. The comparison showed that areas exhibiting decreased backscatter in our SAR-derived map, most notably the shoreline zones of Jamaica Bay and nearby low-lying neighborhoods such as Howard Beach and Spring Creek Park, aligned with FEMA's designated V and A flood risk zones (1% annual-chance flood areas). The correspondence between our observed inundation signatures and these modeled high-risk zones confirms that Sentinel-1 detects water residue in the same vulnerable geographies highlighted by FEMA's long-term hydrological modeling. This comparison supports the credibility of SAR-based flood detection and demonstrates that Sentinel-1 can serve as a near-real-time complement to existing flood risk mapping systems.

This study, while demonstrating that Sentinel-1 can successfully detect floods when properly preprocessed and displayed, did expose weaknesses in Sentinel-1's data collection frequency and how these weaknesses can make Sentinel-1 ineffective in producing enough data for accurate flood prediction. As mentioned before, Sentinel-1 routinely passes by NYC every 6 days, which can make it difficult to obtain data for certain major meteorological events that are not within Sentinel-1's reach of NYC; however, Sentinel-1's 6-day cycle allows relatively consistent observation for a year-long period. The spaced-out periods make Sentinel-1 efficient for observing topographic changes by month. There are geostationary satellites that are consistently pointed at the continental United States, such as NASA's GOES series, which sound accurate in theory until the loss in resolution is considered, and come with a satellite further from the surface of the Earth.

In terms of materials and methodology, we came to realize that QGIS, a program considered crucial for analyzing Sentinel-1 flood maps, was not integral to our final process in analyzing NYC flood data. After encountering error screens and indefinite GeoTIFF loading sequences, our team took a different approach to comparing Sentinel-1 satellite data with actual terrain features seen in NYC to identify high-risk areas. The majority of the visualization took place in Sentinel-1's signature software, SNAP, combined with comparison and analysis using Google Earth Pro, which proved far more efficient than our previous attempts in QGIS. QGIS's role was minimized to simply masking the five boroughs and creating a proper GIS map format of the flood data.

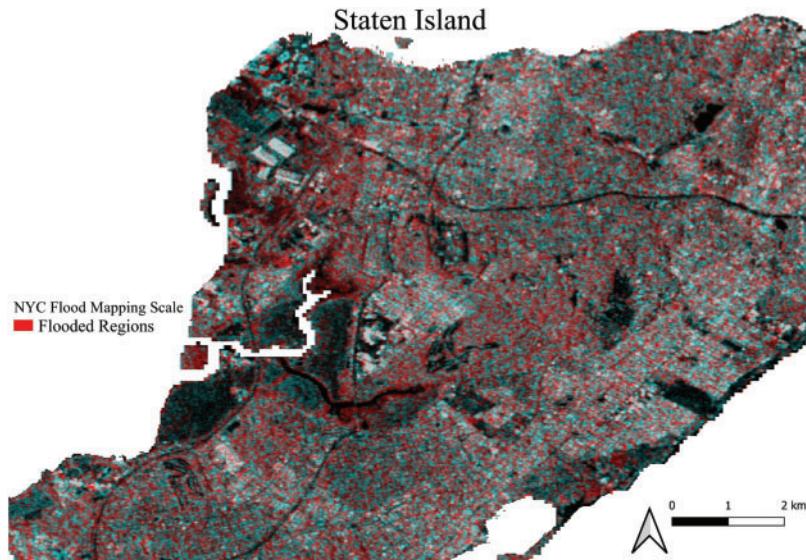


Figure 11. Staten Island on March 24, 2024

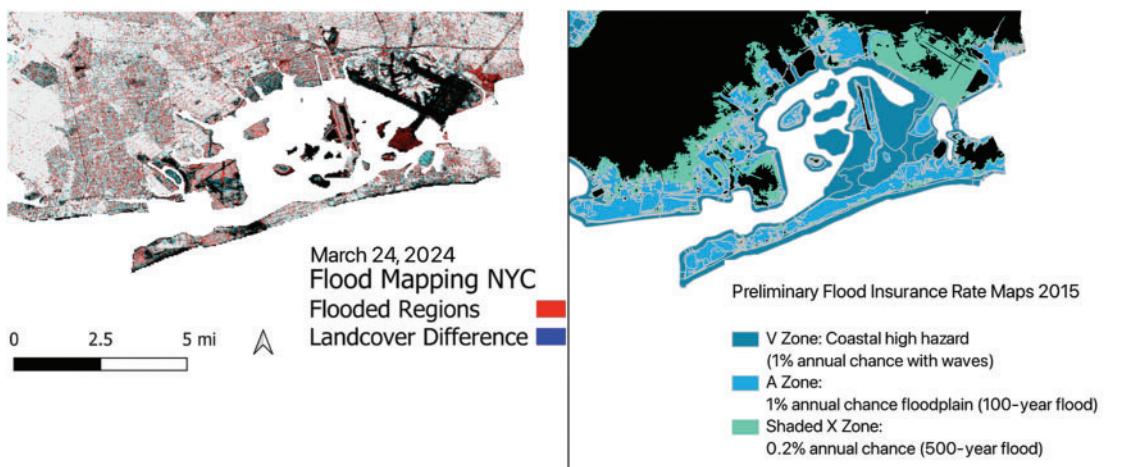


Figure 12. Comparison of Sentinel-1 SAR flood mapping results (March 24, 2024) with FEMA's Preliminary Flood Insurance Rate Map (FIRM) for New York City, obtained from the NY.gov ArcGIS Flood Hazard Viewer

A similar study was conducted in Morocco in 2020 using Sentinel-1 satellite data, further proving the versatility of the satellite regarding its ability to obtain datasets from across the globe.²⁰ Their method for the final analysis to visualize vulnerable flood areas was to use backscatter coefficients of the various Inaouene watershed floods, where lower values are equal to water and high values are equal to non-water areas, to isolate areas that have been flooded. Despite the slightly differing methodology, both processes ultimately lead to similar results where the presence of a flood is proved both visually and analytically.

Conclusion

We obtained data products from Sentinel-1 through ESA's publicly available Copernicus Browser directory. Then, the satellite's personalized SNAP software was used to preprocess the images and highlight the flooded areas for the respective dates on the data products were obtained, using

RGB bands and image stacking. Once fully preprocessed, comparisons were made between the red, high-risk flooded areas that we visualized in Google Earth Pro.

Sentinel-1 has immense potential to provide high-quality data for flood mapping in New York. Future work may involve combining the data from Sentinel-1 with contemporary methods in flood mapping, such as statistical analysis and machine learning, to prevent and minimize the damage of floods. Additionally, optical sensors could be implemented to enhance the temporal resolution of Sentinel-1. Overall, Sentinel-1 serves as a proper foundation, indicating how future satellites can be used to mitigate the impacts of floods.

Acknowledgments

Our deepest gratitude is extended to Prof. Naresh Devineni for his guidance during the project. We also wish to thank the Director of CUNY CREST HIRES, Dr. Shakila Merchant,

and the United Nations University (UNU) Hub on Remote Sensing and Sustainable Innovations for Resilient Urban Systems (R-SIRUS) at The City College of New York for making this research possible and for providing accessible engineering education to students across NYC.

References

- [1] Committee on Urban Flooding in the United States, Program on Risk, Resilience, and Extreme Events, Policy and Global Affairs, Water Science and Technology Board, Division on Earth and Life Studies, and National Academies of Sciences, Engineering, and Medicine. *Framing the Challenge of Urban Flooding in the United States*. Washington D.C: National Academies Press; 2019. doi:10.17226/25381.
- [2] Agonafir C, Lakhankar T, Khanbilvardi R, Krakauer N, Radell D, Devineni N. A review of recent advances in urban flood research. *Water Secur.* 2023;19(2):100141. doi:10.1016/j.wasec.2023.100141.
- [3] Hamlington BD, Bellas-Manley A, Willis JK, et al. The rate of global sea level rise doubled during the past three decades. *Commun Earth Environ.* 2024;5(1):601. doi:10.1038/s43247-024-01761-5.
- [4] Islam T, Zeleke EB, Afroz M, Melesse AM. A systematic review of urban flood susceptibility mapping: remote sensing, machine learning, and other modeling approaches. *Remote Sens.* 2025;17(3):524. doi:10.3390/rs17030524.
- [5] National Sea Level Explorer – U.S. Sea Level Change. 2024. https://earth.gov/sealevel/us/national-sea-level-explorer/?state=NY&scope=section_1.
- [6] NYC Mayor's Office of Climate & Environmental Justice. *Coastal Surge Flooding*. New York City: Climate Impact Spotlight. <https://www.nyc.gov/content/climate/pages/coastal-surge-flooding>.
- [7] New York State Climate Impacts. *Climate Impact Spotlight: New York City*. 2024. <https://nysclimateimpacts.org/explore-by-region/new-york-city/>.
- [8] New York City Department of Environmental Protection. *The New Normal: Combating Storm-Related Extreme Weather in New York City*. 2021. <https://www.nyc.gov/assets/orr/pdf/publications/WeatherReport.pdf>.
- [9] Ten Veldhuis JAE, Clemens FHLR, Sterk G, Berends BR. Microbial risks associated with exposure to pathogens in contaminated urban flood water. *Water Res.* 2010;44(9):2910–2918. doi:10.1016/j.watres.2010.02.009.
- [10] New York City Department of Environmental Protection. *Municipal Separate Storm Sewer System (MS4)*. 2024. <https://www.nyc.gov/site/dep/water/municipal-separate-storm-sewer-system.page>.
- [11] HEAT.gov. July 2025 was planet's 3rd warmest on record. *U.S. Heat.gov*; August 12, 2025. <https://www.heat.gov/pages/urban-heat-islands>.
- [12] New York City Department of Environmental Protection. *Green Infrastructure Rain Gardens*. 2025. <https://www.nyc.gov/site/dep/water/rain-gardens.page>.
- [13] Levy Z, Smardon R, Bays J, Meyer D. A point source of a different color: identifying a gap in United States regulatory policy for “Green” CSO treatment using constructed wetlands. *Sustainability*. 2014;6(5):2392–2412. doi:10.3390/su6052392.
- [14] Rebuild by Design. *Climate Displacement in NYC: Making Space for Our Neighbors*. 2022. <https://rebuildbydesign.org/uncategorized/climate-displacement-in-nyc/>.
- [15] NYC Mayor's Office. Mayor Adams announces next generation of New York City's investments in flood prevention and preparation to keep New Yorkers safe, protect their wallets. *Press Release*; September 9, 2024. <https://www.nyc.gov/mayors-office/news/2024/09/mayor-adams-next-generation-new-york-city-s-investments-flood-prevention-and>.
- [16] Yang W, Xu K, Ma C, et al. A novel multi-objective optimization framework to allocate support funds for flash flood reduction based on multiple vulnerability assessment. *J Hydrol.* 2021;603(3):127144. doi:10.1016/j.jhydrol.2021.127144.
- [17] Torres R, Snoeij P, Geudtner D, et al. GMES Sentinel-1 mission. *Remote Sens Environ.* 2012;120(4):9–24. doi:10.1016/j.rse.2011.05.028.
- [18] De Zan F, Monti Guarnieri A. TOPSAR: terrain observation by progressive scans. *IEEE Trans Geosci Remote Sens.* 2006;44(9):2352–2360. doi:10.1109/TGRS.2006.873853.
- [19] European Space Agency (ESA). *Tracking Maritime Traffic—Sentinel-1 Mission, Copernicus Programme*. 2024. https://www.esa.int/Applications/Observing_the_Earth/Copernicus/Sentinel-1/Tracking_maritime_traffic.
- [20] Benzougagh B, Frison P-L, Meshram SG, et al. Flood mapping using multi-temporal Sentinel-1 SAR images: a case study—Inaouene watershed from Northeast of Morocco. *Arab J Geosci.* 2021;14(2):683. doi:10.1007/s40996-021-00683-y.
- [21] Weather Spark. *March 2024 Weather History at New York City, Central Park, New York, United States*. Weather Spark; 2024. <https://weatherspark.com/h/m/147190/2024/3/Historical-Weather-in-March-2024-at-New-York-City-Central-Park;-New-York;-United-States>.
- [22] Cedar Lake Ventures Inc. *New York City July 2023 Historical Weather Data (New York, United States)*. Weather Spark; September 18, 2025. <https://weatherspark.com/h/m/23912/2023/7/Historical-Weather-in-July-2023-in-New-York-City-New-York-United-States>.
- [23] European Space Agency. *Sentinel Application Platform (SNAP)*. <https://step.esa.int/main/>.
- [24] Cara C. *STEP—Scientific Toolbox Exploitation Platform*. European Space Agency; 2019. <https://step.esa.int/main/>.
- [25] Esri. What is orthorectified imagery? *Esri Insider Blog*. September 18, 2025. <https://www.esri.com/about/newsroom/insider/what-is-orthorectified-imagery>.
- [26] Fresh Kills: Landfill to Landscape. 2022. https://www.nyc.gov/assets/planning/download/pdf/plans/fkl/about_fkl.pdf.
- [27] NYC Parks. *Freshkills Park*. Retrieved 2025. <https://www.nyccgovparks.org/park-features/freshkills-park/about-the-site>.
- [28] University at Albany. *NYS Mesonet: Real-Time and Historical Weather Data*. 2025. <https://www.nysmesonet.org/>.
- [29] Newark Liberty International Airport. *March 2024 Weather History in New York City*. Weather Spark; 2024. <https://weatherspark.com/h/m/23912/2024/3/Historical-Weather-in-March-2024-in-New-York-City-New-York-United-States>.
- [30] Larosa M, Huff J. Saturday soaker gives NYC its third wettest March day ever. *NBC New York*; March 24, 2024. <https://www.nbcnewyork.com/weather/weather-stories/weather-ny-today-forecast-rain-winter-storm/5249411/>.
- [31] Shivonne A. NYC weather: heavy rain causes flight delays, roadway advisories, flooding. *Fox 5 News*; May 23, 2024.

<https://www.fox5ny.com/news/nyc-storm-watch-flights-road-delays-advisory>.

[32] New York State Governor's Office. Governor Hochul updates New Yorkers on state preparations for potential flood impacts as more rain is expected today and tomorrow. *Press Release*; July 16, 2023. <https://www.governor.ny.gov/news/governor-hochul-updates-new-yorkers-state-preparations-potential-flood-impacts-more-rain>.



Solving Problems with Bounds on Linear Forms in Logarithms

Krish Agrawal*

Submitted: 16 September 2025 Accepted: 08 October 2025 Publication date: 29 October 2025

DOI: 10.70671/k3gz8m22

Abstract: This expository paper explores the theory of linear forms in logarithms and its applications to Diophantine equations. We begin with foundational results on transcendental numbers, including Liouville's theorem and the Gelfond–Schneider theorem, before developing Baker's theory of linear forms in logarithms. The paper concludes with applications to Diophantine equations through the Baker–Davenport method, illustrating these techniques with concrete examples. The purpose of this paper is to provide a step-by-step understanding of this particular area of mathematics, and was written as a part of Euler Circle's Independent Paper and Research Writing Program, in which the author studied the topic independently and wrote this paper over a period of 4 weeks.

Author keywords: Linear forms in logarithms; Transcendental number theory; Diophantine equations; Baker's theorem; Number Theory

Introduction

The study of linear forms in logarithms sits at the intersection of number theory and transcendental number theory. It provides tools to solve equations of the form:

$$\alpha_1^{b_1} \alpha_2^{b_2} \cdots \alpha_n^{b_n} = 1$$

by studying the associated linear form:

$$\Lambda = b_1 \log \alpha_1 + \cdots + b_n \log \alpha_n$$

However, before we explore this form of expression further, we first need to understand the basics of transcendental number theory.

Algebraic and Transcendental Numbers

Basic definitions

Definition 2.1 (Algebraic Number). *A complex number α is algebraic if there exists a nonzero polynomial $P \in \mathbb{Z}[x]$ such that $P(\alpha) = 0$. The minimal degree of such a polynomial is called its degree.*

Example 2.2. The following are algebraic:

- $\sqrt{2}$ is algebraic of degree 2, since it satisfies $x^2 - 2 = 0$.
- The golden ratio $\phi = \frac{1 + \sqrt{5}}{2}$ is algebraic of degree 2.
- Any rational number $\frac{p}{q}$ is algebraic of degree 1 via $qx - p = 0$.

It is also helpful to note that algebraic numbers are countable, since they arise as roots of polynomials with integer coefficients.

Definition 2.3 (Transcendental Number). *A complex number is transcendental if it is not algebraic; that is, it does not satisfy any nonzero polynomial equation with integer coefficients.*

Liouville's theorem and constructed transcendents

Joseph Liouville's 1844 theorem was the first major result to establish a criterion for transcendence.

Theorem 2.4 (Liouville's Approximation Theorem). *For any irrational algebraic number α of degree $d \geq 2$, there exists a constant $C(\alpha) > 0$ such that for all rational numbers $\frac{p}{q}$ with $q > 0$:*

$$\left| \alpha - \frac{p}{q} \right| > \frac{C(\alpha)}{q^d}$$

Proof. Let $P(x) = a_d x^d + \cdots + a_0 \in \mathbb{Z}[x]$ be the minimal polynomial of α . By the mean value theorem, for any rational $\frac{p}{q} \neq \alpha$:

$$|P(p/q) - P(\alpha)| = |P'(\xi)| \cdot \left| \alpha - \frac{p}{q} \right|$$

for some ξ between α and p/q . Since P is irreducible, $P(p/q) \neq 0$ and $q^d P(p/q) \in \mathbb{Z}$, so:

$$|P(p/q)| \geq \frac{1}{q^d}$$

*Corresponding Author: Krish Agrawal. Email: krish.595.2026@doonschool.com
The Doon School

Siegel's lemma guarantees bounded p_{km} with:

$$\max |p_{km}| \leq (3N)^{8dN/\log N}$$

Why this bound? The exponential and derivative structure inflates the entries in the system; this bound ensures that a small integer solution still exists.

Refined Schwarz Lemma: If f has n zeros in $|z| < R$, then for $|z| = r < R$,

$$|f(z)| \leq |f(0)| \prod_{k=1}^n \frac{R}{|z_k|} \left(\frac{r}{R}\right)^n \max_{|w|=R} |f(w)|$$

Extending the zeros

Using the identity $\alpha^{z+n} = \alpha^z \cdot \alpha^n$, we obtain:

$$f(\beta + n) = \sum_{k,m} p_{km} \alpha^{k\beta} \alpha^{kn} (\beta + n)^m = \alpha^{n\beta} g(n)$$

where

$$g(z) = \sum_{k,m} p_{km} \alpha^{kz} (z + \beta)^m$$

So $f(\beta + n) = 0$ for $n = 0, 1, \dots, N - 1$.

Order of vanishing

Let s be the smallest integer such that $f^{(s)}(\beta) \neq 0$. Then, $s \geq \tau$ and f have many zeros in a disk around β .

Maximum modulus estimate

We now use a tool from complex analysis to deduce the upper bounds for the function.

Maximum Modulus Principle: If f is analytic in a domain D and continuous on \bar{D} , then $|f(z)|$ attains its maximum on the boundary ∂D .

On the circle $|z| = R = 2N^{1/2}$:

$$|f(z)| \leq (L+1)^2 \cdot \max |p_{km}| \cdot |\alpha|^{LR} \cdot R^L$$

Substituting:

$$|f(z)| \leq N \cdot (3N)^{8dN/\log N} \cdot |\alpha|^{2N} \cdot (2N^{1/2})^{N^{1/2}}$$

Taking logs:

$$\begin{aligned} \log |f(z)| &\leq \frac{8dN}{\log N} \log(3N) + 2N \log |\alpha| + N^{1/2} \log(2N^{1/2}) \\ &\leq CN \end{aligned}$$

Schwarz lemma refinement

To complement the global bound obtained via the maximum modulus principle, we invoke a refined Schwarz-type lemma to estimate the derivatives near β . This allows us to bound $|f^{(s)}(\beta)|$ in terms of the global maximum modulus and the order of vanishing, setting up a contradiction with the algebraic lower bound, which we will see in the subsequent subsection.

We get:

$$\max |f(z)| \leq \left(\frac{R}{N}\right)^{N\cdot\tau} = \left(\frac{2}{N^{1/2}}\right)^{N^2/\log N}$$

So:

$$\log |f(z)| \leq -\frac{N^2}{2} + o(N^2) \Rightarrow |f(z)| \leq e^{-cN^2/\log N}$$

Lower bound via Diophantine approximation

Since α^β is assumed algebraic and $f^{(s)}(\beta) \neq 0$, we apply the following:

Liouville's Theorem (for algebraic numbers): Let ξ be algebraic of degree $d \geq 2$. Then, there exists $C(\xi) > 0$ such that:

$$\left| \xi - \frac{p}{q} \right| > \frac{C(\xi)}{q^d}$$

for all rationals p/q .

We write $f^{(s)}(\beta)$ as a polynomial in α, β , and α^β , with integer coefficients bounded using earlier results. Hence,

$$|f^{(s)}(\beta)| \geq (3N)^{-8d^2N}$$

Why this bound? This bound is derived by estimating the minimal polynomial of the algebraic number $f^{(s)}(\beta)$. The exponent arises from the number of terms and height estimates in the auxiliary function.

Cauchy integral bound and contradiction

Using Cauchy's integral formula:

$$f^{(s)}(\beta) = \frac{s!}{2\pi i} \int_{|z|=R} \frac{f(z)}{(z - \beta)^{s+1}} dz$$

Taking absolute values:

$$|f^{(s)}(\beta)| \leq \frac{s! \cdot 2\pi R}{2\pi (R/2)^{s+1}} \cdot \max |f(z)| = \frac{s! \cdot 2^{s+1} \cdot e^{-cN^2/\log N}}{R^s}$$

Now $s \geq \tau = N/\log N$ implies $s! \leq e^{DN/\log N}$ and $R^s \geq (2N^{1/2})^{N/\log N}$. So we get:

$$|f^{(s)}(\beta)| \leq \frac{e^{DN/\log N} \cdot e^{-cN^2/\log N}}{N^{N/(2\log N)}} = e^{-c'N^2/\log N}$$

But this contradicts the lower bound:

$$(3N)^{-8d^2N} \leq |f^{(s)}(\beta)| \leq e^{-c'N^2/\log N}$$

Taking logs:

$$-8d^2N \log(3N) \leq -\frac{c'N^2}{\log N} \Rightarrow \text{Contradiction for large } N$$

□

Proof framework

The proof of the Gelfond–Schneider theorem reflects a general template used in transcendence theory:

- Construct an auxiliary function with controlled complexity.
- Impose vanishing conditions at many points.
- Use functional equations to extend the zeros.
- Bound the size of the function from above using complex analysis.
- Establish lower bounds via Diophantine approximation.
- Derive a contradiction between the two bounds.

This methodology has been extended to:

- Values of the exponential function at algebraic points (Hermite–Lindemann).
- Logarithms of algebraic numbers (Gelfond–Schneider).
- Elliptic and abelian functions at algebraic points (Schneider).
- Periods of algebraic varieties (conjecturally).

Remark 3.3. *The Gelfond–Schneider theorem resolved part of Hilbert’s 7th problem. It proves that $2^{\sqrt{2}}$ is transcendental, and that $e^{\pi i} = -1$ involves transcendentality as well.*

Baker’s Theory of Linear Forms in Logarithms

²

The Gelfond–Schneider theorem represents a major advance in transcendence theory, settling Hilbert’s 7th problem and providing the first general results for the transcendence of values like α^β , where both α and β are algebraic. However, many problems in number theory involve more complicated expressions—particularly linear combinations of several logarithms of algebraic numbers. A natural question arises: What can be said about expressions such as

$$\Lambda = b_1 \log \alpha_1 + b_2 \log \alpha_2 + \cdots + b_n \log \alpha_n,$$

where α_i are fixed nonzero algebraic numbers and b_i are integers? If such a form is nonzero, how small can it be in terms of the sizes of the coefficients and the complexity of the numbers involved?

In the 1960s, Alan Baker developed a powerful generalization of the Gelfond–Schneider theorem, proving that any nontrivial linear form in logarithms of algebraic numbers is not only nonzero but also bounded away from zero by an explicit, effective lower bound. This result, now known as Baker’s theorem, laid the foundation for a large part of effective Diophantine analysis and provided new methods for bounding the size of solutions to exponential equations.

We now state a simplified version of Baker’s theorem, emphasizing its effectiveness and applicability.

Theorem 4.1 (Baker, 1966). *Let $\alpha_1, \dots, \alpha_n$ be nonzero algebraic numbers, none equal to 1, and suppose that the logarithms $\log \alpha_i$ are taken with respect to a fixed branch of the logarithm on $\mathbb{C} \setminus (-\infty, 0]$. Let $b_1, \dots, b_n \in \mathbb{Z}$, not all zero. Define*

$$\Lambda = b_1 \log \alpha_1 + \cdots + b_n \log \alpha_n.$$

Then, if $\Lambda \neq 0$, we have

$$|\Lambda| > \exp(-C_1 \cdot \log B \cdot \log A_1 \cdots \log A_n),$$

where

- $B = \max\{|b_1|, \dots, |b_n|\}$,
- $each A_i \geq \max\{Dh(\alpha_i), |\log \alpha_i|, 0.16\}$,
- $h(\alpha_i)$ is the (logarithmic) height of α_i ,
- D is the degree of the number field $\mathbb{Q}(\alpha_1, \dots, \alpha_n)$,
- and C_1 is an effectively computable constant depending only on n and D .

This bound is significant for several reasons. First, it gives a concrete inequality—if one knows the coefficients b_i and the heights and degrees of the α_i , then the right-hand side becomes a computable quantity. Second, the bound allows one to conclude that Λ is not too small. In many number-theoretic applications, this means that the expression Λ must actually be bounded away from zero, and therefore an integer linear combination of logarithms cannot vanish.

The logarithmic height $h(\alpha)$ measures the arithmetic complexity of an algebraic number. For rational numbers $\alpha = \frac{p}{q}$ in lowest terms, we have $h(\alpha) = \log \max\{|p|, |q|\}$. For general algebraic numbers, this extends via consideration of the minimal polynomial and the archimedean absolute values of its conjugates. The appearance of $h(\alpha_i)$ in the lower bound reflects the intuitive idea that the more complicated the input numbers are, the more flexibility they allow in small linear combinations.

For a full and rigorous proof of Baker’s theorem, the reader is referred to *Transcendental Number Theory* by Alan Baker.²

It is also worth emphasizing that Baker’s theorem is effective: the constant C_1 can be computed explicitly given the input data. This is a major departure from earlier transcendence theorems, which often proved existence results without providing any computational information.

The power of this theorem is best appreciated through its applications. For instance, it enables one to solve Diophantine equations involving exponential expressions by first proving that a certain linear form in logarithms is nonzero and then bounding how small it can be. These bounds, though typically extremely small, are nonetheless finite and lead directly to constraints on the size of possible integer solutions.

In the next section, we will see how this theoretical result can be made more practical using a method developed by Baker and Davenport, which allows one to reduce the resulting bounds significantly by incorporating ideas from Diophantine approximation and continued fractions.

Computing Effective Bounds Using Baker's Theorem

The true power of Baker's theorem lies not just in its qualitative assertion of transcendence or linear independence, but in its capacity to provide explicit lower bounds on expressions involving logarithms of algebraic numbers. This is crucial in many Diophantine contexts, where the goal is to prove that certain equations admit only finitely many solutions, or even to identify those solutions completely.

Suppose we have algebraic numbers $\alpha_1, \dots, \alpha_n$, each not equal to 0 or 1, and define the linear form:

$$\Lambda = b_1 \log \alpha_1 + \dots + b_n \log \alpha_n$$

where the coefficients b_1, \dots, b_n are integers, not all simultaneously zero. We assume $\Lambda \neq 0$. The question is: How small can $|\Lambda|$ be?

Baker's theorem provides a compelling answer. If we let $B = \max |b_i|$, let $h(\alpha_i)$ denote the logarithmic height of α_i , and let D be the degree of the number field generated by all α_i , then there exists a computable constant C_1 , depending only on n and D , such that:

$$|\Lambda| > \exp(-C_1 \cdot (\log B)(\log A_1) \cdots (\log A_n))$$

Here, each A_i is defined as:

$$A_i = \max\{Dh(\alpha_i), |\log \alpha_i|, 0.16\}$$

To see this in action, consider the expression $\Lambda = x \log 2 - y \log 3$, which arises in problems like finding solutions to $2^x = 3^y$. If $x, y \leq 1000$, then $B = 1000$, and both 2 and 3 are rational, so $h(2) = \log 2$, $h(3) = \log 3$, and $D = 1$.

Plugging in, the bound becomes:

$$|\Lambda| > \exp(-C_1 \cdot \log(1000) \cdot \log \log 2 \cdot \log \log 3)$$

With standard estimates for the logarithms and C_1 on the order of 10^4 , we find that the lower bound is something like 10^{-10000} . While seemingly negligible, this value is not zero—and that makes all the difference. This bound guarantees a minimum separation between linear combinations of logarithms of algebraic numbers, allowing us to eliminate hypothetical integer solutions with large coordinates.

Although the exponential decay in the bound may feel extreme, especially for large B , its explicit nature makes it useful in practice. It opens the door to computational refinements such as the Baker–Davenport method, which we now explore.

Despite the strength of Baker's bounds, their numerical size often makes them impractical for direct computation. The Baker–Davenport method addresses this limitation. It combines the theoretical guarantees from Baker's theorem with elementary tools like continued fractions to dramatically reduce the bound and isolate the few possible solutions. This hybrid approach bridges the gap between transcendental number theory and concrete problem-solving.

Matveev's Refinement of Linear Form Bounds

Baker's theorem provides an effective lower bound for linear forms in logarithms, but the constants involved can be extremely large in practice, limiting the usability of the bounds in computational settings. In 2000, E. M. Matveev introduced a major refinement of these bounds, leading to significantly sharper estimates and better constants for a wide class of problems.³

Matveev's theorem applies to linear forms in logarithms of algebraic numbers and provides fully explicit lower bounds that are often several orders of magnitude smaller than those obtained from Baker's original result. These improvements are especially noticeable in applications involving small-degree algebraic numbers or when the logarithmic heights of the numbers are moderate.

In simplified form, Matveev's bound states that for algebraic numbers $\alpha_1, \dots, \alpha_n \in \overline{\mathbb{Q}}^\times$, and nonzero integers b_1, \dots, b_n , if $\Lambda = b_1 \log \alpha_1 + \dots + b_n \log \alpha_n \neq 0$, then

$$|\Lambda| > \exp(-C(n) \cdot D^2 \cdot (1 + \log D)(1 + \log B) \cdot (\log A_1) \cdots (\log A_n)),$$

where

- D is the degree of the number field generated by the α_i ,
- $B = \max\{|b_i|\}$,
- $A_i \geq h(\alpha_i)$ are bounds on the logarithmic heights,
- and $C(n)$ is an explicit constant depending only on n .

Although this expression resembles Baker's original bound structurally, Matveev's constants are much more favorable. This makes it possible to use the resulting inequalities directly in computations, particularly in bounding solutions to exponential Diophantine equations.

Modern applications often rely on Matveev's version of the bound for practical results. For example, in determining perfect powers in linear recurrence sequences or resolving variants of the Ramanujan–Nagell equation, Matveev's bound can reduce the upper bounds for unknowns from astronomical values to within computational reach.

Remark 6.1. For a full proof of Matveev's result and its many variants, the reader is referred to his 2000 paper in the *Izvestiya: Mathematics journal*.³

The Baker–Davenport Reduction Method

While Baker's theorem provides an explicit lower bound on $|\Lambda|$, the bound is often so tiny that it alone cannot exclude large integer solutions from consideration. The Baker–Davenport method bridges this gap. It strengthens the application of Baker's bound by combining it with elementary yet powerful tools from Diophantine approximation, particularly continued fractions. This method was introduced in a collaborative paper by Baker and Davenport.⁴

The basic idea is to recast the inequality $|\Lambda| > \delta$ into a statement about how well a rational number can approximate a certain irrational one. Suppose we have:

$$\Lambda = x \log \alpha - y \log \beta \neq 0$$

We divide both sides by $x \log \beta$ to obtain:

$$\left| \frac{\log \alpha}{\log \beta} - \frac{y}{x} \right| > \frac{\delta}{|x \log \beta|}$$

This inequality shows that $\frac{y}{x}$ is not too close to the irrational number $\log \alpha / \log \beta$. If we can find rational approximations to $\log \alpha / \log \beta$ that are closer than the bound allows, we can rule them out as candidates. Conversely, we may use the convergents of the continued fraction expansion of $\log \alpha / \log \beta$ to approximate this ratio and compare them to known integer solutions.

This technique allows us to zoom in on potential solutions with small values of x and y , sharply reducing the upper bounds that arise from Baker's theorem. In this way, what was initially a massive bound, say $x < 10^{20}$, can often be brought down to something like $x < 100$, making a brute-force search feasible.

Thus, the Baker–Davenport method transforms a theoretical lower bound into a practical algorithm. It is this combination of transcendental number theory and computational approximation that makes the method one of the most effective tools in solving exponential Diophantine equations.

Worked example: solving $2^x - 3^y = 1$

Consider the equation $2^x - 3^y = 1$. This problem asks for pairs of integers (x, y) such that the difference between a power of two and a power of three equals one. Rearranging, we get:

$$2^x = 3^y + 1$$

Taking logarithms, we find:

$$x \log 2 = \log(3^y + 1)$$

For large y , $3^y + 1$ is very close to 3^y , and so:

$$x \log 2 \approx y \log 3$$

which leads to:

$$\left| \frac{\log 2}{\log 3} - \frac{y}{x} \right| \approx \text{very small}$$

Now, using the continued fraction expansion of $\log 2 / \log 3 \approx 0.6309$, we compute convergents such as $2/3$, $3/5$, $8/13$, and so on. Each of these gives a candidate rational approximation y/x . For each such pair (x, y) , we check whether it satisfies the original equation.

Trying $x = 1$, we find:

$$\begin{aligned} 2^1 &= 2, \quad 3^0 = 1, \quad 2 - 1 = 1 \Rightarrow \text{solution found: } (x, y) \\ &= (1, 0) \end{aligned}$$

Trying $x = 2$, we find:

$$\begin{aligned} 2^2 &= 4, \quad 3^1 = 3, \quad 4 - 3 = 1 \Rightarrow \text{solution found: } (x, y) \\ &= (2, 1) \end{aligned}$$

Trying $x = 4$ yields $2^4 = 16$, but $3^2 = 9$ and $16 - 9 = 7$, which does not satisfy the equation. As we continue this process, we find that all other convergents violate the bound imposed by Baker's theorem and refined by the continued fraction approximation.

Eventually, this method eliminates all possible solutions except $(1, 0)$ and $(2, 1)$. In this way, the Baker–Davenport method provides a complete and rigorous resolution of the Diophantine equation.

Applications to Diophantine Equations

Perfect powers in recurrence sequences

One particularly striking application of linear forms in logarithms arises in the analysis of perfect powers appearing within classical recurrence sequences. A natural example is the Fibonacci sequence (F_n) , defined by $F_0 = 0$, $F_1 = 1$, and the recurrence relation $F_{n+2} = F_{n+1} + F_n$.

The question we now pose is whether the Fibonacci sequence contains any perfect squares beyond the trivial examples. Indeed, it is easy to check that $F_0 = 0$, $F_1 = F_2 = 1$, and $F_{12} = 144 = 12^2$. But are there any others?

To address this, one begins by recalling the closed-form expression for F_n , known as Binet's formula:

$$F_n = \frac{\phi^n - \psi^n}{\sqrt{5}}$$

where $\phi = \frac{1 + \sqrt{5}}{2}$ is the golden ratio and $\psi = \frac{1 - \sqrt{5}}{2}$ is its algebraic conjugate. As n grows, the term ψ^n becomes negligible since $|\psi| < 1$, so we approximate:

$$F_n \approx \frac{\phi^n}{\sqrt{5}}$$

Now, suppose F_n is a perfect square, say $F_n = y^2$ for some integer y . Then:

$$\frac{\phi^n}{\sqrt{5}} \approx y^2 \Rightarrow \phi^n \approx \sqrt{5}y^2$$

Taking logarithms on both sides, we obtain:

$$n \log \phi \approx \log \sqrt{5} + 2 \log y$$

and hence,

$$|n \log \phi - 2 \log y - \log \sqrt{5}| \approx 0$$

This is now a linear form in logarithms of algebraic numbers: $n \log \phi - 2 \log y - \log \sqrt{5}$. Applying Baker's theorem to this expression provides an effective lower bound on its absolute value, which in turn leads to an explicit upper bound on n . Typically, the bound from Baker's theorem is large, but it can be substantially reduced using the Baker–Davenport method.

Once an upper bound is known, one simply checks all Fibonacci numbers F_n for n up to that bound and verifies whether any of them is a perfect square. Such an analysis shows that the only values of n for which F_n is a perfect square are $n = 0, 1, 2, 12$. No other Fibonacci numbers are perfect squares.

This type of argument, which combines the closed-form representation of a recurrence with bounds on linear forms in logarithms, has proven extremely effective in answering questions about perfect powers in recurrence sequences more generally.

The Thue equation

Another important application of Baker's theory involves solving Thue equations. These are Diophantine equations of the form:

$$F(x, y) = m$$

where $F(x, y)$ is an irreducible homogeneous binary form of degree at least three with integer coefficients, and m is a nonzero integer. The classic result due to Thue tells us that for a fixed F and m , this equation has only finitely many integer solutions (x, y) . However, Thue's original proof is ineffective—it gives no means to compute or bound the solutions.

Baker's theory changes this entirely. It allows us to transform the problem into a system involving linear forms in logarithms, thereby furnishing explicit upper bounds on $|x|$ and $|y|$.

The argument begins by factorizing the binary form $F(x, y)$ over the algebraic closure of \mathbb{Q} . Suppose we have:

$$F(x, y) = \prod_{i=1}^d (\alpha_i x - \beta_i y)$$

where the α_i, β_i lie in some finite extension of \mathbb{Q} , and $d \geq 3$ is the degree of F . Now, since $F(x, y) = m$, we see that for at least one index i , the quantity $|\alpha_i x - \beta_i y|$ must be small—no larger than roughly $|m|^{1/d}$.

This gives us an approximation:

$$\left| \alpha_i \frac{x}{y} - \beta_i \right| \ll |y|^{-d}$$

Taking logarithms, we consider expressions of the form:

$$\log \left| \alpha_i \frac{x}{y} - \beta_i \right|$$

These are precisely the kinds of quantities to which Baker's bounds apply. Using the theory of linear forms in logarithms, we can place explicit lower bounds on such logarithmic expressions. Matching this with the upper bounds obtained from the factorization and the size of m , we deduce explicit upper bounds on $|x|$ and $|y|$.

Once those bounds are known, even if large, they reduce the original infinite problem to a finite one: we simply need to check all integer pairs (x, y) within the bounded region to find all solutions to the Thue equation.

In practice, further refinements, often using reduction techniques and continued fraction approximations, can decrease the computational load significantly. This makes Baker's theory not only theoretically satisfying but also practically viable for determining the complete set of solutions to many Diophantine equations once thought intractable.

Conclusion

Linear forms in logarithms form a central part of modern transcendental number theory, providing explicit techniques for dealing with exponential Diophantine equations. What distinguishes this area is not only its ability to prove the transcendence of certain numbers but also its effectiveness in yielding concrete numerical bounds on the size of integer solutions. This paper has shown how classical theorems like those of Liouville and Gelfond–Schneider lay the foundation for Baker's general result. Through Baker's theorem and its computational refinements, especially the Baker–Davenport method, we obtained tools capable of producing meaningful results in practice. The examples examined, from exponential equations such as $2^x = 3^y + 1$ to the study of perfect powers in recurrence sequences, illustrate how linear forms in logarithms can be used to fully resolve equations that otherwise admit infinitely many possibilities.

Although the bounds provided by Baker's theorem are often too large for direct application, methods like continued fraction approximation allow one to reduce them substantially. This combination of transcendental estimates with classical number-theoretic techniques is what gives the method its strength.

Acknowledgments

I would like to express my sincerest gratitude to my mentors, Justin Wu and Simon Rubinstein-Salzedo, without whose guidance and support I would not have been able to explore the field of transcendental number theory or pursue this particular topic.

References

- [1] Gelfond AO. *Transcendental and Algebraic Numbers*. Dover Publications; 1960.
- [2] Baker A. *Transcendental Number Theory*. Cambridge University Press; 1975.
- [3] Matveev EM. An explicit lower bound for a homogeneous rational linear form in logarithms of algebraic numbers. *Izv Math*. 2000;64(6):1217–1269.
- [4] Baker A, Davenport H. The equations $3x^2 - 2 = y^2$ and $8x^2 - 7 = z^2$. *Q J Math*. Oxford Series; 1969;20(2):129–137.
- [5] Lang S. *Introduction to Transcendental Numbers*. Addison-Wesley; 1966.
- [6] Shorey TN, Tijdeman R. Exponential diophantine equations. In: *Cambridge Tracts in Mathematics*. Cambridge University Press; 1986.

About the Authors



Krish is currently a high school senior studying at The Doon School pursuing the IB Diploma, and intends to study Physics in University. His paper is an expository paper exploring Linear Forms in Logarithms, an important result in the field of transcendental number theory.



Firm Size and Change in Employment over the Economic Cycle

Mahi Malhotra*

Submitted: 16 August 2025 Accepted: 14 October 2025 Publication date: 29 October 2025

DOI: [10.70671/y50y7n25](https://doi.org/10.70671/y50y7n25)

Abstract: The objective of this research is to understand whether, during economic downturns (and upturns) in the United States, employment in small businesses is disproportionately more affected than that in large firms. During the coronavirus disease 2019 pandemic, small firms struggled to recover from their losses and therefore had to cut back on their number of employees despite efforts from federal and state governments to provide emergency relief to small businesses. For a firm size that employs almost half of the private sector employees, the results of this research are critical in identifying solutions to alleviate the burden on businesses disproportionately affected during recessions.

This paper examines the relationship between the change in employment in small versus large firms as compared to the annual real gross domestic product (GDP) growth rate over the 1979 to 2021 time period. Results show that the percentage change in employment for small firms does indeed show a stronger linear relationship with the real GDP growth rate than the corresponding relationship for large firms. Expanding fiscal support measures to help keep small firms afloat during downturns would indirectly influence their employment and reduce the decline in employment. Additionally, the Small Business Administration should prioritize loans primarily based on financial need to combat this disparity and spread resources in a way that benefits more small businesses.

Author keywords: Firm size; Unemployment; Economic Cycle

Introduction

Small businesses are one of the most influential drivers of growth in the economy.¹ Despite their name, small businesses employ 46.4% of private sector employees and have been the main contributors to job creation for the past few years.² However, small and local businesses have often been portrayed as more vulnerable than large firms to the fluctuations in the economic cycle. If the employment of small firms is indeed particularly volatile, then this makes a large proportion of the workers in the economy quite vulnerable during recessions. On the other hand, if results indicate that it is actually the employment of large businesses that is cyclically volatile, then efforts should be made to lessen the impact recessions have on the employment of large firms.

This paper examines the percentage change in employment in small and large firms from 1979 to 2021 and the annual real gross domestic product (GDP) growth rate. Percentage change in employment was used to measure employment growth, as this variable would specifically track year-to-year differences in the percentage of employees lost or gained, providing a direct insight into the proportion of employees leaving or entering a particular firm size. Sahin et al.³ also used the percentage change in employment as

their variable of interest in a similar study. Measuring the percentage change in employment allows for proportional comparison between firm sizes. Results show that the percentage change in employment for small firms does indeed show a stronger linear relationship with the real GDP growth rate than the percentage change in employment for large firms.

Placing this research in the broader literature, there generally seems to be more studies supporting the idea that employment in large firms varies with the economic cycle more than in small firms.⁴⁻⁶ However, during the coronavirus disease 2019 (COVID-19) pandemic, small firms were affected much more significantly than large firms.⁷ This is not reflected in the existing literature, as most of it is pre-pandemic; these studies lack more recent data. Additionally, this paper and other literature vary in terms of the type of metric used to assess employment. For example, while this study utilized annual percentage change in employment for both firm sizes, other studies used annual net job creation. Finally, different studies define “small” and “large” firms in different ways. While a small business is defined as one having fewer than 500 employees and a large firm as one having 500 or more employees, which aligns with the Small Business Administration (SBA) definition, other studies may refer to small businesses as having fewer than 50 employees and large firms as having more than 1,000 employees. This disparity could certainly lead to different results and conclusions.

*Corresponding Author: Mahi Malhotra. Email: mahi.malhotra819@gmail.com
Millburn High School, Millburn, NJ 07041

This paper is organized as follows: The Literature Review section covers a deeper analysis of past research on employment across firm sizes during economic cycles. The Description section presents the datasets used, the variables analyzed, and the methodology of the data analysis. Afterward, two ordinary least-squares regression lines (OLSRLs) are introduced: one for the percentage change in employment for small firms and one for the percentage change in employment for large firms. The OLSRLs, which show similar correlations, are compared based on the coefficients of the explanatory variables of interest, as well as any other statistically significant variables. The Discussion section explores potential solutions to combat the dip in employment of small businesses during recessions. Finally, the Conclusion section takes note of the purpose and significance of the results.

Literature Review

Recovery in employment after a recession is slower now than it was in the past. Recessions prior to 1990 generally saw a steady recovery in employment soon after the recessions ended.⁸ Furthermore, recent recessions until the Great Recession show a continued decline in employment even during economic recoveries. Additionally, the continued decline in employment could be correlated with a rise in the number of firms of a particular size class, such as the recent and rapid rise in the number of small businesses in the United States since the 1990s.² Though a link between the two has not been proven, it is certainly possible and therefore important to consider while reviewing older literature.

Some papers provide evidence supporting the sensitivity of small firms to economic fluctuations. The recessions in the early 1990s and 1980s saw the net job creation rates of small firms suffer more than those of large firms.⁵ Additionally, during the COVID-19 pandemic, fiscal support measures putting small firms at a disadvantage were prominent, and policies to aid small businesses were insufficient. While large businesses had multiple avenues to seek relief, small businesses could only be aided by the Paycheck Protection Program and Health Care Enhancement Act (PPP). Additionally, the priority of handing out PPP loans was given to more established businesses, so the most in-need small businesses struggled.⁹ Furthermore, data from the recession of 2007–2009 indicate that small firms were hit harder than large firms because of reduced customer demand (especially in small firms), and therefore a decline in fixed investment by small firms.³ The research presented in this paper focuses more on the relationship between recessions and employment than others. It is worth mentioning that GDP growth was found to have a stronger correlation with the employment growth at smaller firms than at larger ones.⁵

However, not everyone agrees with the idea that small firms are more cyclically sensitive than large firms. It is still an ongoing debate. Other evidence points to large firms as being more sensitive to economic fluctuations.^{3–5} Regarding net employment growth rates, large firms suffered more than small firms during the recessions of 1982 and 2001. Even

in the Great Recession, the growth rate of employment at initially large firms declined by 1.65% more.⁵ Additionally, in the recession of 2001, the percentage change in employment of large firms was more greatly affected than that of small firms.³ In terms of net job creation rate, large businesses proportionately shed more jobs over recessions and create more jobs during expansions.⁴ There appears to be a pattern that relatively poor job creation of large firms happens for years after the trough. Large firms are typically more productive and offer higher-paying jobs, so they can successfully poach workers from smaller competitors. Job-to-job reallocation redistributes workers from low- to high-paying jobs.¹⁰ This research often links lower-paying firms to small firms and higher-paying firms to large firms, which may not always be an accurate representation of large and small firms. During downturns, large firms want to get rid of the excess of workers they obtained from poaching.⁵ This indicates that large firms will proportionately shrink more than small firms during economic contractions. For example, Moscarini and Postel-Vinay claim that small, credit-constrained firms would quickly rebound during recessions due to the central bank's reactions in monetary policy.⁵ As the economy recovers and unemployment falls, monetary policy tightens and curbs the growth of small firms, which are outperformed by larger, financially less constrained competitors, again causing the larger firms to proportionately gain more employees.

Historically, research from 1994 indicates that in the prior six recessions, only the 1970 and the 1980 recessions saw a clear collapse in the growth rate of sales at small firms in comparison to large firms.⁶ While there is a clear collapse of the growth rate of sales of small firms relative to large firms in the recessions of 1970 and 1980, the opposite occurs in the recession of 1961, and the recessions of 1974, 1982, and 1991 appear fairly neutral. The poor performance of small businesses during both recessions lasted for just a year, while in all recoveries, large firms remained sluggish for years.

Building upon data from the pandemic, new research disputes some ideas presented in older studies.⁹ Previous studies have linked higher wages to larger firms, but this correlation has weakened in recent decades. The authors found that during the COVID-19 pandemic, there was greater excess unemployment among workers in small enterprises than in larger firms. However, in industries with higher flexibility for working remotely, the pattern reversed, and small businesses retained their workers at slightly higher rates than larger firms. It is important to note that this reversed pattern is only seen for micro-businesses with fewer than 10 employees, which are likely family businesses or ones with close ties.

Small businesses were more likely to be in the nonessential services sector, so they were particularly affected during the pandemic. Workers in small- and medium-sized enterprises are more likely to have lower education levels and live in rural areas, contributing to the firm size wage premium, although since 1980, the firm size wage premium has declined substantially, especially for low-skilled workers.⁹ Still, there is significant variability in the wages of workers in small firms.

Recent years have seen increased product and labor market concentration in the economy due to scale-based

technological changes and barriers to the entry of new competitors, leading to the rise of a few superstar firms. The measure of new COVID-19 cases per 1,000 residents was also used to show that small firms faced a surge in unemployment compared to large firms.⁹ For every additional case per 1,000 residents, unemployment increased by more than 2.2% for firms with fewer than 10 workers, 2% for firms with 10–99 residents, 1.6% for firms with 100–999 workers, and 1.2% for firms with more than 999 employees. The authors highlighted that the pandemic augmented existing inequalities between small and large firms, as there was a further increase in product and labor market concentration as more small businesses failed and cut back their workforces.

Description

Data

For this research, data were sourced from the U.S. Census Bureau's Business Dynamic Statistics (BDS), Federal Reserve Economic Data (FRED), and the United Nations Department of Economic and Social Affairs Population Division. The data source for the statistics used to calculate the annual percentage change in employment for large and small firms was the BDS. The BDS is a set of datasets that provide measures of business dynamics (such as job creation and destruction, establishment births and deaths, and firm startups and shutdowns) for the economy overall and aggregated by establishment and firm characteristics. The datasets are based on administrative and survey-collected data. The BDS contains data at the establishment level, although the data analyzed in this research were at the parent firm level. The dataset utilized, "Sector by Firm Age by Initial Firm Size," contains annual statistics on the number of employees by firm size (1–19 employees, 20–499 employees, and 500 or more employees). In the context of this paper, small firms are defined as firms with 1 to 500 employees, while large firms are defined as firms with 500 or more employees. This definition of "small" and "large" aligns with the SBA's definitions. As per the BDS, firm size is the average number of employees in the current and prior years.

Additionally, raw data on real GDP as well as the annual unemployment rate were derived from FRED, which is an online database of time series regarding different economic variables based on U.S. data. It is compiled from data provided by the U.S. Census Bureau as well as the Bureau of

Labor Statistics, and is managed by the Federal Reserve Bank of St. Louis.

Finally, the crude rate of net migration in the United States, defined as the ratio of net migration during the year to the average population in that year and expressed per 1000 persons, is derived from the UN's Crude Rate of Net Migration data. This data is obtained from the United Nations Department of Economic and Social Affairs Population Division. The datasets used to calculate the crude rate of net migration include the U.S.'s Total Net-Migration and the U.S.'s Total Population by Sex.

Variables

The explanatory variables in this analysis include the real GDP annual growth rate, real GDP annual growth rate with a lag of 1 year, U.S. net migration growth rate, unemployment rate, and the percentage change in employment by firm size (the data for small firm sizes were used for the analysis of small firms, and the data for large firm sizes were used for the analysis of large firms) with a lag of 1 year. These variables were considered to be potentially correlated with the percentage change in employment for both small and large firms, as the crude rate of net migration affects the labor force and the unemployment rate would likely align with changes in employment for the two firm sizes. The response variable is the percentage change in employment by firm size (either small or large). [Table 1](#) presents summary statistics for the variables used in the analysis.

Methodology

The annual percentage change in employment between 1979 and 2021 was calculated as shown in [Eq. \(1\)](#).

$$\% \text{ change in number of employees} = 100 * (\text{number of employees in the current year} - \text{number of employees in the prior year}) / (\text{number of employees in the prior year}). \quad (1)$$

To plot against the annual percentage change in employment for large and small firms, the growth rate of real GDP was calculated as shown in [Eq. \(2\)](#).

$$\text{growth rate of real GDP} = 100 * (\text{real GDP of the current year} - \text{real GDP of the prior year}) / (\text{real GDP of the prior year}). \quad (2)$$

Table 1. Summary statistics

Response variables	Summary statistics				
	Q1	Median	Q3	Mean	Std. Dev.
Percentage change in employment for small firms	-0.284	1.285	2.354	0.948	2.486
Percentage change in employment for large firms	0.969	2.100	3.322	1.656	2.235
Real GDP annual growth rate	1.915	2.768	3.829	2.622	1.985
Crude rate net migration	2.651	4.495	4.942	4.021	1.276

This data was calculated beginning in 1979 as well. More explanatory variables were added to the equation as controls due to possible correlations they may have with the response variable. The crude rate of net migration in the United States was calculated using $1000 * (\text{total net migration in the United States for a given year} / \text{total population of the United States for that given year})$ from 1979 to 2021. The annual unemployment rate from FRED was kept as is. Additionally, a lag(1) variable of the annual percentage change in employment was added as an explanatory variable because the previous year's percentage change in employment may relate to the next year's. Furthermore, a lag(1) variable of the annual growth rate of real GDP was added as well. Due to the lag of 1 year, all the data are measured from 1980 to 2021.

The following regression model was considered to investigate the effect on percentage change in employment for the two firm sizes:

$$y_t = \alpha + \beta_1 y_{t-1} + \beta_2 x_t + \beta_3 x_{t-1} + \beta_4 m_t + \beta_5 u_t + e_t, \quad (3)$$

where y_t is the annual percentage change in employment (for small or large firms) at time t ; α is the model intercept; y_{t-1} is the annual percentage change in employment at $t - 1$, included to account for persistence in the variable of interest. The main coefficients of interest are β_2 , which captures the sensitivity of the percentage change in employment to the annual growth rate of GDP at time t (x_t), and β_3 , which accounts for the delayed effect. Finally, the model incorporates m_t , the crude rate of net migration at time t , and u_t , the annual unemployment rate at time t . Parameters β_4 and β_5 are regression coefficients corresponding to m_t and u_t , respectively. The term e_t is the error term. The regression parameters in Eq. (3) were estimated by performing an OLS regression in Microsoft Excel.

Results

Overall, the OLS estimate for the relationship between the real GDP annual growth rate at time t and the percentage change in employment at time t was 0.342 (Table 2) for large firms and 0.667 (Table 3) for small firms. The value of 0.667 is notably larger than 0.342, indicating that the percentage change in employment for small firms is much more strongly correlated with the real GDP annual growth rate than the percentage change in employment for large firms is. These

results suggest a stronger sensitivity to the economic cycle for employment at small firms. The least squares estimate for the relationship between the lagged real GDP annual growth rate and the percentage change in employment further supports this conclusion, given that the coefficient is 1.110 for small firms (Table 3) and 0.828 (Table 2) for large firms. All these values are statistically different from 0.

Additionally, only for the data for small firms, the coefficient of the percentage change in employment for small firms with a lag of 1 year is statistically significant at -0.306 (Table 3). Since the coefficient of the percentage change in employment lag(1) is only significant for the OLSRL of small firms, it implies that the prior year's percentage change in employment correlates with the current year's percentage change in employment far more for small firms than for large firms. Since this is a negative slope, the percentage change in employment lag(1) and the percentage change in employment have an inverse relationship for the OLSRL of small firms. Perhaps there is a yearly cyclability in regard to the percentage change in employment for small firms, alternating between growth and decline each year.

All other coefficients of variables, including the percentage change in employment for large firms lag(1), the crude rate of net migration, and the unemployment rate, are not statistically significantly different from zero (Tables 2 and 3). It is possible that any observed association between these variables and the percentage change in employment for small and large firms occurred simply by chance. However, it is also possible that this study may be underpowered, and that more data pre-1979 would reveal coefficients for each variable that are statistically significant. The BDS's earliest data regarding the number of employees based on firm size were from 1978, so this study was limited to calculating the percentage change in the number of employees from 1979 to 2021 (when the data ended).

The two OLSRLs calculated for their respective firm sizes have reasonably strong positive R^2 values: 0.705 for the OLSRL of small firms (Table 5) and 0.741 for the OLSRL of large firms (Table 4). It is interesting to note that the standard error of the OLSRL of small firms, 1.441 (Table 5), is somewhat greater than the standard error of the OLSRL of large firms, 1.214 (Table 4).

Table 2. Regression coefficients for large firms

	Coeff.	Std. err.	t stat	P-value	Lower 95%	Upper 95%
Intercept	-3.076	1.970	-1.561	0.127	-7.072	0.920
Percentage change in employment for large firms lag(1)	0.167	0.126	1.331	0.192	-0.088	0.422
Real GDP annual growth rate	0.342	0.119	2.870	0.007	0.100	0.584
Real GDP annual growth rate lag(1)	0.828	0.120	6.889	0.000	0.584	1.071
Crude rate net migration	0.268	0.176	1.523	0.136	-0.089	0.624
Unemployment rate	0.052	0.175	0.295	0.770	-0.304	0.408

Table 3. Regression coefficients for small firms

	<i>Coeff.</i>	<i>Std. err.</i>	<i>t stat</i>	<i>P-value</i>	<i>Lower 95%</i>	<i>Upper 95%</i>
Intercept	−5.446	1.959	−2.780	0.009	−9.420	−1.473
Percentage change in employment for small firms lag(1)	−0.306	0.145	−2.102	0.043	−0.600	−0.011
Real GDP annual growth rate	0.667	0.126	5.286	0.000	0.411	0.923
Real GDP annual growth rate lag(1)	1.110	0.176	6.319	0.000	0.754	1.467
Crude rate net migration	0.019	0.208	0.092	0.928	−0.403	0.441
Unemployment rate	0.337	0.179	1.888	0.067	−0.025	0.700

Table 4. Regression statistics for large firms

Regression statistics	
Multiple R	0.861
R ²	0.741
Adjusted R ²	0.705
Standard error	1.214
Observations	42

Table 5. Regression statistics for small firms

Regression statistics	
Multiple R	0.840
R ²	0.705
Adjusted R ²	0.664
Standard error	1.441
Observations	42

Discussion

Results of this study indicate that small firms are more sensitive to the economic cycle. Many research papers have contributed to the discussion of the effects of size class on employment over the economic cycle, and different metrics have been used to measure employment. For example, Moscarini and Postel-Vinay⁴ used the net job creation rate. They identified that gross job creation for large firms surges ahead of that of small firms at the peak of expansions. The paper recognizes this same pattern occurring within industries as opposed to across industries, although the data it uses extend only up to 2006. It has also been observed that the negative correlation between the net job creation rate of large employers and the level of unemployment at business cycle frequencies is stronger for larger employers than for smaller employers. Moscarini and Postel-Vinay⁵ studied net employment growth rates, finding that large firms suffered more than small firms during the recessions of 1982 and 2001. In addition to different metrics being used, research papers also differ in their definitions of “small” and “large” businesses. For example, Sahin et al.³ used the percentage

change in employment between small (1–49 employees), medium (50–499 employees), and large (500+ employees) firms as a metric to measure employment and directly compare data from the Great Recession and the dot-com crash.

Based on the analysis presented in this paper, the percentage change in employment of small firms has a stronger relationship with the annual growth rate of GDP, as well as with the annual growth rate of GDP with a lag of 1 year. The slope for the annual growth rate of GDP lag(1) for small firms is slightly steeper than 1, indicating that this variable is a particularly good predictor of the percentage change in employment of small firms.

Protecting the employment security of almost half of the private sector’s employees during economic downturns should be the SBA’s priority. Using the coefficients from the OLSRL in this paper for the annual growth rate of GDP, as well as (more importantly) the annual growth rate of GDP lag(1), the SBA can predict when small businesses may see a decline in employment. Small firms often rely heavily on external bank loans and credit. SBA loans are known for having low interest rates, which is ideal for many small businesses, but it is tough to qualify for them. Often-times, more well-established businesses get priority when it comes to loans that the SBA hands out, so there are still many younger, more vulnerable businesses struggling to stay afloat. The SBA should prioritize loans primarily based on financial need to combat this disparity and spread resources in a way that benefits more.

Conclusions

This paper investigated the effects of the economic cycle on the percentage change in employment for small and large firms. Key sources of data included employee data from the U.S. Census Bureau’s BDS and annual GDP data from the FRED. A regression model was used to investigate the effect on the percentage change in employment for the two firm sizes. The results of this research demonstrate that small firms are more sensitive to the economic cycle than large firms, which contrasts with the results of several studies in this area.

Being able to learn and apply the coefficient of the OLSRL for the annual growth rate of GDP lag(1) for small firms to predict the percentage change in employment for

the next year is particularly vital. While it may not be able to foresee completely unexpected events, such as the COVID-19 pandemic, it can certainly be of aid during economic contractions. Above all else, this paper should bring more attention to employees of small businesses and the effort to find ways to protect their jobs, whether this includes reforming how the SBA administers loans or providing small businesses with other avenues to obtain credit.

With different definitions of small and large firms throughout the long discourse on which firm size is more cyclically variable, this research used the SBA's definition of a small firm, which is fewer than 500 employees, while a large firm is 500 or more employees. The results of this research contribute to the discussion of employment variability over the economic cycle by firm size by attempting to standardize the definitions of large and small firms. As more research is conducted in this area, having standard definitions of large and small firms would provide a clearer picture of the dynamics between firm size and changes in employment.

Working Paper No. 20884). *Nat Bureau Econ Res.* 2015. doi:10.3386/w20884.

Acknowledgment

I extend my appreciation to Helena Palma Carvalho, who assisted me with processing and analyzing my data, as well as structuring my paper. Additionally, I thank my family for supporting and encouraging me.

References

- [1] Hasan R, Chy MAR, Johora FT, Ullah MW, Saju MAB. Driving growth: the integral role of small businesses in the U.S. economic landscape. *American J Indus Business Manag.* 2024;14(6):256–272. doi:10.4236/ajibm.2024.146043.
- [2] Office of Advocacy. *Frequently Asked Questions About Small Business.* U.S. Small Business Administration; 2024. https://advocacy.sba.gov/wp-content/uploads/2024/12/Frequently-Asked-Questions-About-Small-Business_2024-508.pdf.
- [3] Sahin A, Kitao S, Cororaton A, Laiu S. Why small businesses were hit harder by the recent recession. *SSRN Elect J.* 2011;17(4). doi:10.2139/ssrn.1895527.
- [4] Moscarini G, Postel-Vinay F. Large employers are more cyclically sensitive (NBER Working Paper No. 14740). *Nat Bureau Econ Res.* 2009. doi:10.3386/w14740.
- [5] Moscarini G, Postel-Vinay F. The contribution of large and small employers to job creation in times of high and low unemployment. *American Econ Rev.* 2012;102(6):2509–2539. doi:10.1257/aer.102.6.2509.
- [6] Gertler M, Gilchrist S. Monetary policy, business cycles, and the behavior of small manufacturing firms. *Quart J Econ.* 1994;109(2):309–340. doi:10.2307/2118465.
- [7] *Did the Pandemic Hit Small Firms Harder than Large Firms?* (n.d.). www.stlouisfed.org/on-the-economy/2022/aug/pandemic-hit-small-firms-harder-than-large-firms.
- [8] Gordon RJ. Okun's law and productivity innovations. *American Econ Rev.* 2010;100(2):11–15. doi:10.1257/aer.100.2.11.
- [9] Lin K-H, Aragão C, Dominguez G. Firm size and employment during the pandemic. *SAGE Open.* 2021;11(2): 2378023121992601. doi:10.1177/2378023121992601.
- [10] Moscarini G, Postel-Vinay F. Cyclical reallocation of workers across employers by firm size and firm wage (NBER



Hybrid Deep Learning Approaches for Classifying Autism from Brain MRI

Ashley Chen*

Submitted: 17 October 2025 Accepted: 29 October 2025 Publication date: 29 October 2025

DOI: 10.70671/mdqw4m29

Abstract: Background: Autism spectrum disorder (ASD) is most often diagnosed using behavioral evaluations, which can vary between clinicians. Brain imaging, combined with machine learning, may help identify more objective patterns linked to ASD. In this research, magnetic resonance imaging (MRI) data from the publicly available Autism Brain Imaging Data Exchange I (ABIDE I) dataset ($n = 1,112$) were used to test two approaches for classifying individuals with ASD and typically developing control participants. Both groups were drawn from the same ABIDE I dataset, which includes MRI scans from individuals diagnosed with ASD as well as neurotypical individuals without any major neurological or psychiatric conditions. The first approach used a three-dimensional convolutional neural network (CNN) trained end-to-end, while the second employed a hybrid method in which the CNN served as a feature extractor followed by a support vector machine (SVM) classifier. Results showed that the baseline CNN achieved moderate performance, whereas the hybrid CNN + SVM model demonstrated higher overall accuracy and produced more balanced results between ASD and control groups. These findings indicate that separating feature extraction from classification can improve performance and reduce bias between diagnostic categories. Overall, this study suggests that deep learning methods may enhance the reliability and objectivity of MRI-based research on ASD.

Author keywords: Autism spectrum disorder; magnetic resonance imaging; MRI; convolutional neural networks; support vector machines; biomarkers

Introduction

Autism spectrum disorder (ASD) is a neurodevelopmental condition that affects social communication, behavior, and sensory processing.¹ Diagnosis currently relies on behavioral assessments such as the Autism Diagnostic Observation Schedule 2 (ADOS-2) and the Autism Diagnostic Interview-Revised (ADI-R). While these tools are widely used, they are usually subjective and may depend on individual interpretation. As a result, diagnostic outcomes may vary across evaluators and settings.²

The need for more objective diagnostic methods has grown alongside rising prevalence rates. According to data from the U.S. Centers for Disease Control and Prevention (CDC), approximately one in 150 children were identified with ASD in 2000, compared to one in 36 by 2020 (Table 1). This increasing prevalence highlights the importance of developing tools that can support early and consistent identification.

Neuroimaging offers one possible path toward objective biomarkers. Advances in machine learning, particularly deep learning methods such as convolutional neural networks (CNNs), offer new opportunities to capture subtle

and spatially distributed neuroanatomical differences that may not be evident with traditional analysis. CNNs have shown strong performance in medical imaging, including mammography,⁴ brain tumor grading,⁵ and diabetic retinopathy screening.⁶ These successes suggest that CNNs are well suited for ASD research, where subtle structural variations across multiple brain regions must be identified.

Despite their potential, CNNs face several challenges. They require large datasets to achieve reliable generalization,⁷ yet many neuroimaging studies in ASD involve only a few hundred participants. Models are also highly sensitive to variability in scanner hardware and acquisition protocols. These differences can produce site-specific artifacts that may be mistaken for disorder-related features.⁸ In addition, CNNs often function as “black boxes,” with limited interpretability even when tools such as saliency maps or class activation methods are applied.⁹ These limitations complicate efforts to build clinically robust biomarkers and highlight the importance of evaluating model generalizability.

This study investigates whether CNNs applied to structural MRI can distinguish individuals with ASD from typically developing controls (TDCs). A baseline end-to-end CNN was compared to a hybrid framework in which CNN-derived features were classified with support vector machines (SVM). The aim was to determine whether separating feature learning from classification improves accuracy and

*Corresponding Author: Ashley Chen.

Email: ashleylalachen@gmail.com

Oakton High School, 2900 Sutton Rd, Vienna, VA 22181, USA

Table 1. Rates of ASD diagnosis³

Surveillance year	Birth year	Number of sites reporting	Combined prevalence per 1,000 children (Range across sites)	This is about 1 in X children
2020	2012	11	27.6 (23.1–44.9)	1 in 36
2018	2010	11	23.0 (16.5–38.9)	1 in 44
2016	2008	11	18.5 (18.0–19.1)	1 in 54
2014	2006	11	16.8 (13.1–29.3)	1 in 59
2012	2004	11	14.5 (8.2–24.6)	1 in 69
2010	2002	11	14.7 (5.7–21.9)	1 in 68
2008	2000	14	11.3 (4.8–21.2)	1 in 88
2006	1998	11	9.0 (4.2–12.1)	1 in 110
2004	1996	8	8.0 (4.6–9.8)	1 in 125
2002	1994	14	6.6 (3.3–10.6)	1 in 150
2000	1992	6	6.7 (4.5–9.9)	1 in 150

Table 2. Demographic and clinical characteristics of participants across ABIDE I sites¹⁰

Site	N	% ASD	% Male	Age (Mean ± Std)
California Institute of Technology	23	47.8	78.3	27.1 ± 5.8
Carnegie Mellon University	32	40.6	78.1	26.8 ± 9.8
Kennedy Krieger Institute	44	25.0	77.3	10.1 ± 1.2
University of Leuven	62	46.8	88.7	18.1 ± 5.0
Ludwig Maximilians University Munich	34	5.9	88.2	25.3 ± 10.3
NYU Langone Medical Center	122	43.4	73.0	13.8 ± 5.8
Olin, Institute of Living at Hartford Hospital	36	55.6	86.1	16.8 ± 3.5
Social Brain Lab BCN NIC UMC Groningen and Netherlands Institute for Neurosciences	17	11.8	100	32.7 ± 7.0
San Diego State University	24	12.5	70.8	14.1 ± 1.9
Trinity Centre for Health Sciences	35	28.6	100	16.8 ± 3.5
University of California, Los Angeles	102	54.9	88.2	13.1 ± 2.5
University of Michigan	129	41.1	81.4	14.2 ± 3.3
University of Pittsburgh School of Medicine	56	51.8	85.7	18.8 ± 6.9
University of Utah School of Medicine	100	57.0	100	22.1 ± 7.7
Yale Child Study Center	34	17.7	70.6	13.1 ± 2.8

Note: Values are reported as sample size (N), percentage diagnosed with ASD, percentage male, and mean age ± standard deviation.

generalization, thereby advancing progress toward objective neuroimaging biomarkers for ASD.

Methodology

Dataset

This study used structural magnetic resonance imaging (sMRI) data from the Autism Brain Imaging Data Exchange I (ABIDE I), a publicly available repository aggregating neuroimaging and clinical data from 17 international research sites, including both individuals diagnosed with ASD and TDCs. TDCs are participants who do not have ASD or

any other major neurological, developmental, or psychiatric conditions. They are considered neurotypical and serve as the comparison group for evaluating differences relative to individuals with ASD.

In total, 1,112 participants were included, comprising 539 individuals with ASD and 573 controls. Participant ages ranged from childhood to adulthood (mean site ages between ~10 and ~33 years), and the sample reflected the known male predominance in ASD. Table 2 summarizes the distribution of participants across contributing sites, including sample sizes, sex ratios, diagnostic proportions, and mean ages.

Structural MRI provides high-resolution anatomical images of the brain, capturing cortical thickness, white matter, and subcortical morphology. This modality was selected because it offers stable, reproducible measures of brain anatomy suitable for machine learning pipelines.

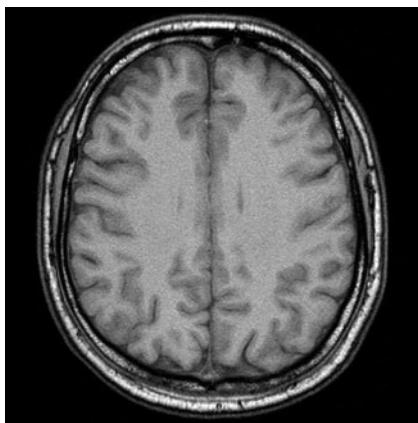


Figure 1. Example of a T1-weighted structural magnetic resonance imaging scan from a typically developing control participant in the Autism Brain Imaging Data Exchange I dataset¹¹

Preprocessing

Before analysis, raw MRI scans underwent extensive preprocessing to account for heterogeneity across the 17 acquisition sites in ABIDE I and to prepare the data for input into CNNs. These steps were essential for reducing scanner-related variability while preserving biologically meaningful structural features relevant to autism. The preprocessing pipeline included the following stages:

- 1. Skull Stripping:** Non-brain tissue such as scalp, skull, and dura was removed using the HD-BET algorithm,¹² a deep-learning based brain extraction tool. This ensured that subsequent analyses focused exclusively on brain tissue.
- 2. Spatial Normalization:** Each scan was registered to the MN152 standard brain template using Advanced Normalization Tools. This step placed all participant images into a common stereotactic space, aligning anatomical structures across subjects and reducing variation due to head position or scanner orientation.
- 3. Intensity Normalization:** MRI intensity values vary systematically across scanners and acquisition protocols, which can introduce site-specific artifacts into multi-site datasets such as ABIDE. To mitigate these effects, voxel intensities were normalized using a **histogram standardization**. This method aligns the intensity distribution of each subject's scan to a common reference distribution, thereby reducing scanner-related variability while preserving biologically meaningful signals.

Formally, each voxel intensity I was transformed as:

$$I' = F_{ref}^{-1}(F(I)) \quad (1)$$

where $F(I)$ is the cumulative distribution function (CDF) of the subject's intensity histogram, and F_{ref}^{-1} is the inverse CDF of a chosen reference histogram (in this case, derived from the median intensity distribution across all participants). This mapping ensures that intensities are rescaled consistently across sites, preventing the CNN from inadvertently learning scanner differences instead of ASD-related features.

- 4. Resampling and Cropping:** Following normalization, volumes were resampled from their original native resolution (approximately $256 \times 256 \times 160$ voxels at 1 mm isotropic spacing, varying by site) to a uniform resolution of $96 \times 96 \times 96$ voxels. This downsampling provided standardized inputs for the CNN, balancing computational efficiency with preservation of fine-grained neuroanatomical detail. Cropping around the brain minimized empty space and further reduced computational overhead.
- 5. Quality Control:** Following preprocessing, all scans were visually inspected to confirm successful skull stripping, registration, and normalization. Scans that failed preprocessing were excluded from further analysis.

In addition to the individual preprocessing steps, an overall schematic of the pipeline was constructed to illustrate the sequential workflow from raw MRI scans to the finalized dataset. [Fig. 5](#) highlights the integration of skull stripping,

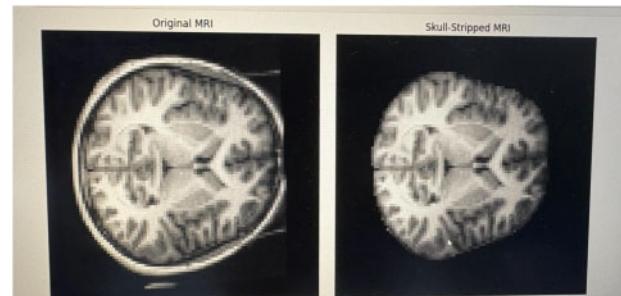


Figure 2. Comparison between an original magnetic resonance imaging scan (left) and the same scan after skull stripping (right), showing removal of non-brain tissue

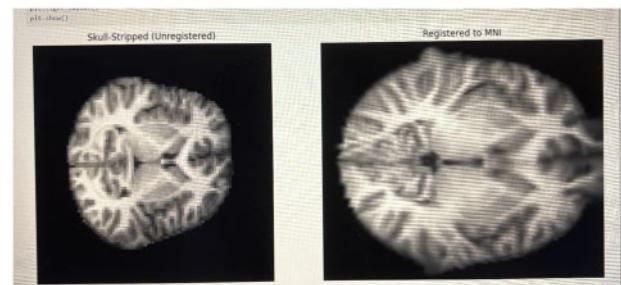


Figure 3. Magnetic resonance imaging scan after skull stripping (left) compared with the same scan registered to the MNI template (right)

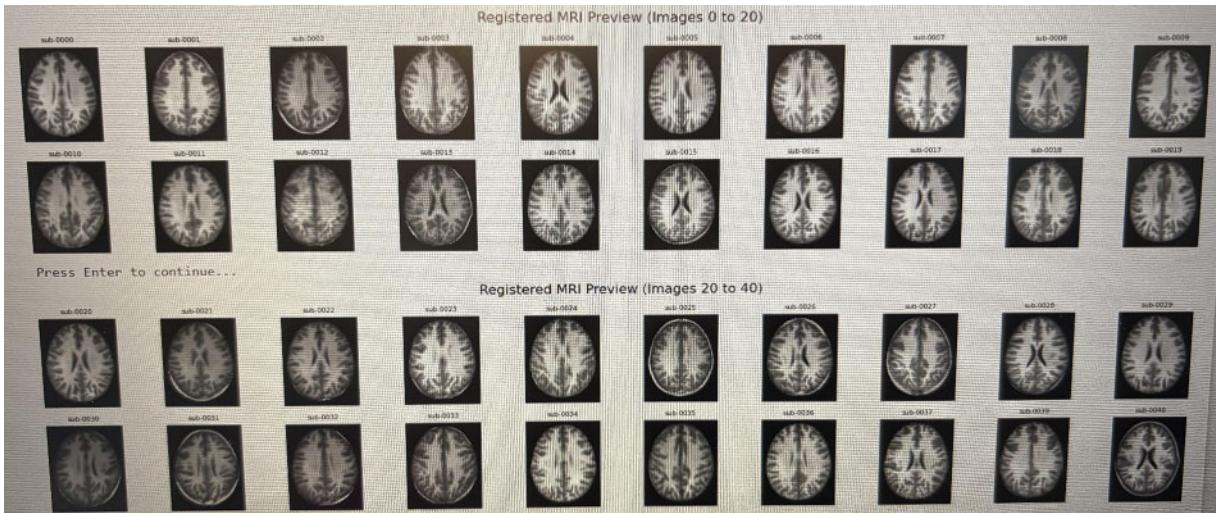


Figure 4. Example of registered magnetic resonance imaging scans across multiple participants after preprocessing. All scans were visually inspected to confirm preprocessing quality

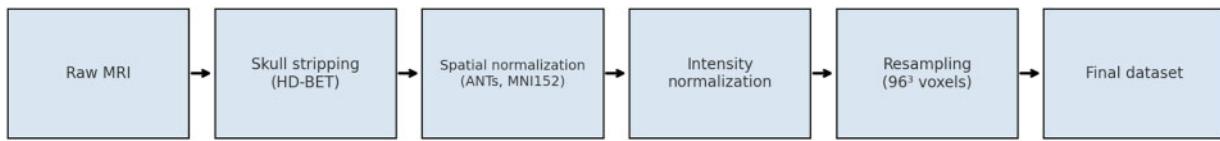


Figure 5. Preprocessing pipeline for structural magnetic resonance imaging scans

spatial normalization, intensity normalization, and resampling into a standardized process that ensured comparability across participants and sites.

Model architectures

Two classification approaches were implemented to investigate the utility of CNNs for distinguishing individuals with ASD from TDC. The first approach consisted of an end-to-end 3D CNN trained directly on preprocessed structural MRI scans. The second approach adopted a hybrid framework, in which CNNs served as feature extractors and the derived representations were classified using an SVM.

Baseline End-to-End CNN. The baseline model was a three-dimensional (3D) CNN designed to learn spatially distributed features across the whole brain. Each input volume had dimensions of $96 \times 96 \times 96$ voxels after preprocessing.

The architecture comprised multiple convolutional layers that scanned the MRI volumes with small filters to detect local structural patterns. Each convolutional operation was followed by a rectified linear unit (ReLU) activation, introducing nonlinearity and enabling the network to capture more complex features. Max-pooling layers were interleaved to downsample the feature maps, reducing dimensionality while retaining salient information.

The final convolutional output was flattened into a 1D feature vector and passed through fully connected layers, which combined information across the brain. The last stage of the network performed classification using a softmax function, which converted the outputs into probabilities for ASD versus control.

Formally, the classifier first computed a pair of unnormalized values, one for ASD and one for control, through a linear transformation:

$$z = Wx + b \quad (2)$$

where x is the feature vector, W and b are learned parameters, and z represents unnormalized class scores. These values, also known as logits, represent unnormalized class scores that can be interpreted as the model's internal evidence for ASD versus control. At this stage, the scores are not probabilities and only become meaningful once transformed by the softmax function:

$$\hat{y}_i = \frac{e^{z^1}}{\sum_{j=1}^2 e^{z^j}}, \quad i \in \{1, 2\} \quad (3)$$

where y is the true diagnostic label (ASD or control). During training, the model compared its predicted probabilities with the true labels and adjusted its parameters to reduce errors. This learning process was optimized using the Adam algorithm. To improve generalization and prevent the model from fitting noise in the data, dropout was applied to randomly deactivate some connections during training, and batch normalization was used to stabilize and speed up learning.

The complete architecture of the baseline 3D CNN is summarized in Fig. 6, which highlights the sequential flow of operations from raw MRI input to diagnostic prediction.

Hybrid CNN + SVM. To improve generalization, the CNN was restructured to act as a feature extractor rather than a full classifier. After convolution, activation (ReLU), and pooling, the final convolutional output was flattened into a one-dimensional feature vector. Instead of passing

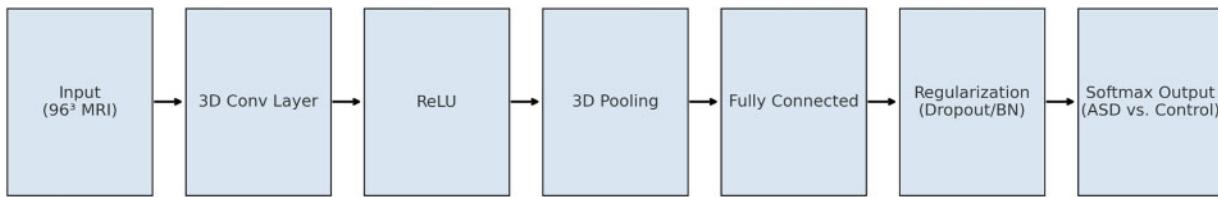


Figure 6. General architecture of the baseline three-dimensional convolutional neural network

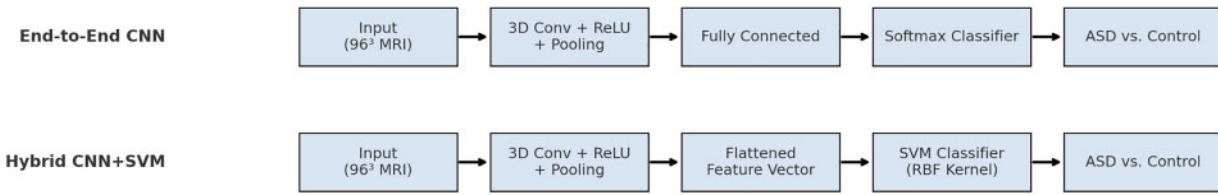


Figure 7. Comparison of end-to-end convolutional neural network (CNN) and hybrid CNN + SVM architectures

this vector through dense layers and a softmax classifier, it was output to an SVM. The rationale was that the CNN could capture complex anatomical features from the MRI scans, while the SVM could provide a more robust separation between ASD and control participants in the resulting feature space.

Formally, the SVM aimed to find a decision boundary that maximized the margin between classes. This was achieved by solving the following optimization problem:

$$\min_{w, b, \xi} \frac{1}{2} \|w\|^2 + C \sum_{i=1}^n \xi_i \quad (4)$$

subject to

$$y_i(w \cdot x_i + b) \geq 1 - \xi_i, \xi_i \geq 0 \quad (5)$$

where x_i is the CNN-derived feature vector for subject i , $y_i \in \{-1, +1\}$ is the class label (control or ASD), w and b define the separating hyperplane, and ξ_i are slack variables that allow some misclassifications. The regularization parameter C controls the trade-off between maximizing the margin and penalizing errors. A radial basis function (RBF) kernel was employed to allow the SVM to capture nonlinear class boundaries in the high-dimensional feature space.

The overall differences between the two classification strategies are illustrated in Fig. 7. In the end-to-end CNN, MRI volumes are processed through convolutional layers, fully connected layers, and a softmax classifier to produce diagnostic probabilities. In contrast, the hybrid CNN + SVM model uses the CNN only to extract features, which are then classified by an SVM with an RBF kernel. This separation of feature learning from classification highlights the alternative pathways by which MRI data can be mapped to diagnostic predictions.

Data partitioning and leakage control

To ensure fair evaluation and to mitigate the influence of site-specific artifacts, data partitioning was performed using a controlled and reproducible strategy. All 1,112 participants were divided into training, validation, and test subsets into a 70:15:15 ratio. Because the ABIDE I dataset combines

scans acquired from 17 independent research sites, each with unique scanner models, acquisition parameters, and population demographics, subjects originating from the same site were kept within the same split. This site-based grouping prevented the model from exploiting scanner-dependent intensity patterns or geometric characteristics, which could otherwise lead to artificially inflated performance.

With each site-based partition, stratification was performed according to the diagnostic label to preserve the ratio of individuals with ASD and TDCs across all subsets. This ensured that each data split reflected the global class distribution of the entire cohort. The partitioning process was implemented through a group-aware splitting procedure that used the site identifier as the grouping variable.

Nested cross-validation with five outer folds was applied to the training data to provide unbiased model selection and performance estimation. The inner loop of this procedure was used to tune hyperparameters such as learning rate, weight decay, and SVM kernel parameters, while the outer loop evaluated generalization to unseen data. This approach reduced the likelihood of overfitting and provided a more stable estimate of expected model performance on independent samples.

All preprocessing operations that required data-dependent fitting, including histogram standardization, intensity normalization, and feature scaling, were restricted to the training data. Parameters derived from the training set were subsequently applied to the validation and test sets to prevent any inadvertent data leakage. To ensure reproducibility, the random seed was fixed to the same value across all computational environments, including Python, NumPy, TensorFlow, and scikit-learn. This ensured that data partitioning and model initialization were identical across repeated experiments.

This partitioning framework provided a rigorous structure for evaluation that minimized data leakage and enhanced reproducibility. The resulting methodology ensured that the trained models were evaluated under conditions that accurately reflected their generalizability to unseen sites and populations.

Class imbalance handling

The ABIDE I dataset contains an unequal number of individuals diagnosed with ASD and TDCs, reflecting the higher prevalence of TDC participants across several contributing sites. This imbalance can bias the model toward predicting the majority class and may result in artificially elevated accuracy without a corresponding improvement in sensitivity to ASD. To mitigate this effect, the training procedure incorporated multiple strategies that adjusted for class frequency while preserving the natural diagnostic proportions in the validation and test sets.

For the CNN, class imbalance was addressed by introducing class-weighted loss optimization. The binary cross-entropy loss function was weighted inversely to class frequency so that errors on ASD samples contributed more strongly to the total loss than errors on TDC samples. The class weight for each label was computed using the following formula:

$$W_c = \frac{N}{2N_c} \quad (6)$$

where N is the total number of samples in the training set and N_c is the number of samples belonging to class c . This weighting ensured that both classes exerted equal influence on parameter updates during backpropagation, despite unequal sample counts. The weighting scheme was applied only during the training phase and was not extended to the validation or test subsets.

For the hybrid CNN + SVM framework, the extracted feature vectors were classified using an SVM trained with the parameter setting `class_weight = "balanced"`, which automatically adjusts the penalty parameter C for each class in proportion to the inverse of its frequency. This adjustment prevents the decision boundary from being dominated by the majority class and promotes balanced classification performance across the ASD and TDC groups. During hyperparameter tuning, the weighting was maintained constant across all inner-fold splits to ensure comparability of performance metrics.

No oversampling, undersampling, or synthetic sample generation was performed, as such methods can distort the underlying neuroanatomical distribution and introduce unrealistic structural variability. Instead, the chosen weighting approach preserved the integrity of the original data while compensating for unequal representation through algorithmic adjustments in the loss and margin functions. This method has been shown in prior neuroimaging studies to maintain biological validity while improving sensitivity to underrepresented diagnostic categories.

Through these weighting procedures, the impact of class imbalance was effectively reduced, allowing the models to focus on meaningful neuroanatomical differences associated with ASD rather than on frequency-driven biases.

Data augmentation

To improve model generalization and reduce overfitting, controlled data augmentation techniques were applied during training. Augmentation introduces controlled variability into the training samples, allowing the model to become

more robust to minor spatial and intensity variations that naturally occur across different MRI acquisitions. In the context of neuroimaging, however, augmentation must be applied conservatively to avoid generating anatomically implausible representations of brain structure.

Each training volume underwent random spatial transformations that preserved the overall geometry of the brain while introducing slight variability in orientation and scale. Specifically, 3D rotations were sampled uniformly within $\pm 7.5^\circ$ around each anatomical axis, and isotropic scaling factors were randomly drawn from the range of 0.95 to 1.05. These transformations simulate small positional differences that can arise from head alignment or scanner calibration without distorting cortical or subcortical morphology. All augmented images were generated on-the-fly during training, ensuring that no identical input was seen twice by the network and that storage requirements remained manageable.

No left-right flipping or nonlinear deformations were applied, as these operations could disrupt hemispheric asymmetries that are biologically meaningful in autism-related neuroanatomical studies. Similarly, intensity augmentations such as histogram perturbations or Gaussian noise injection were avoided to prevent confounding effects on voxel-level contrast, which may represent diagnostically relevant features.

All augmentation procedures were confined strictly to the training subset and were not applied to the validation or test data. This restriction maintained a clear separation between model optimization and evaluation, ensuring that reported performance metrics reflected genuine generalization rather than adaptation to artificial data variability. The augmentation pipeline was implemented using TensorFlow's built-in 3D transformation utilities, with parameters verified through visual inspection to confirm anatomical plausibility.

These carefully controlled augmentation procedures ensured that the model was exposed to realistic spatial variability, enhancing robustness to inter-site differences and improving stability during training while maintaining the integrity of neuroanatomical structures.

CNN training details

The baseline 3D CNN described in Section 2.3 was trained to distinguish individuals with ASD from TDCs using the preprocessed MRI volumes as input. The training procedure was designed to optimize classification performance while minimizing overfitting and ensuring stable convergence across multiple folds of cross-validation.

Training was conducted using the Adam optimization algorithm, which adaptively adjusts learning rates based on estimates of the first and second moments of the gradients. The initial learning rate was set to 1×10^{-4} , and the exponential decay rates for the first and second moment estimates were set to 0.9 and 0.999, respectively. To improve convergence and prevent stagnation at local minima, a learning rate reduction strategy was employed. The learning rate was halved if the validation loss failed to improve after five

consecutive epochs, with a minimum learning rate threshold of 1×10^{-6} .

The model was trained with a batch size of 8 for a maximum of 100 epochs per fold. An early stopping mechanism was implemented to terminate training when the validation loss did not improve for 10 consecutive epochs, preventing unnecessary computation and reducing the risk of overfitting. The model state corresponding to the epoch with the highest validation area under the receiver-operating characteristic curve (ROC-AUC) was preserved for evaluation.

To improve model regularization, several additional techniques were incorporated. A dropout rate of 0.3 was applied to the fully connected layer, randomly deactivating a fraction of neurons during training to promote redundancy and prevent co-adaptation of features. L2 weight regularization with a coefficient of 1×10^{-5} was applied to the convolutional kernels to penalize overly large weights and encourage smoother representations. Batch normalization was inserted after each convolutional layer to stabilize gradient propagation, accelerate convergence, and reduce internal covariate shift. The activation function used throughout the network was the ReLU, chosen for its computational efficiency and ability to mitigate vanishing gradients.

The model was implemented in TensorFlow using the Keras high-level API. Training was conducted on a workstation equipped with an NVIDIA RTX 3080 GPU (10 GB VRAM) and 32 GB of system memory. The computational environment was configured with fixed random seeds across Python, NumPy, TensorFlow, and scikit-learn to ensure full reproducibility. Training progress was monitored using TensorBoard to visualize learning curves, loss trajectories, and AUC performance across epochs.

This training configuration achieved a balance between computational efficiency and model generalization. The combination of adaptive learning rates, regularization techniques, and early stopping contributed to stable convergence across folds while preserving sensitivity to diagnostically relevant neuroanatomical patterns in the structural MRI data.

Hybrid SVM training and tuning

In the hybrid classification framework, the CNN described in Section 2.3 served as a fixed feature extractor, and an SVM was subsequently trained to perform the final diagnostic classification. This approach was designed to leverage the representational strength of deep learning while incorporating the interpretability and stability of traditional machine learning classifiers. By separating feature extraction from classification, the hybrid pipeline allowed the CNN to capture high-dimensional neuroanatomical patterns, while the SVM focused on optimizing class separation within the resulting feature space.

The SVM classifier was trained using an RBF kernel, which projects data into a higher-dimensional space to capture nonlinear relationships between features. A grid search was employed to identify the optimal hyperparameters for the regularization parameter C and the kernel coefficient γ . Specifically, C was explored in the set {0.1, 1, 10, 100}, and γ was explored in the set $\{1 \times 10^{-4}, 1 \times 10^{-3}, 1 \times 10^{-2}, 1 \times 10^{-1}\}$. The grid search was performed within the inner loop of the nested cross-validation framework described in Section 2.4, using validation AUC as the selection criterion. The parameter combination yielding the highest mean validation AUC was retained for the final model evaluation on the outer folds.

To account for the class imbalance present in the dataset, the SVM was trained with `class_weight = "balanced"`, which automatically adjusts the penalty parameter for each class in inverse proportion to its frequency. This configuration ensured that both diagnostic categories contributed equally to the optimization of the separating hyperplane. The optimization was performed using the `libsvm` backend in `scikit-learn`, with a maximum of 10,000 iterations per model to guarantee convergence.

Following training, the learned decision function was applied to the test set to generate continuous decision scores. These scores were later converted to probabilistic estimates using Platt scaling, providing interpretable confidence values for each prediction. The combination of CNN-based feature extraction and SVM-based classification yielded a robust and interpretable hybrid framework that balanced nonlinear representational power with well-calibrated decision boundaries.

Experimental Evaluation

Evaluation metrics and procedure

Model performance was assessed using the nested cross-validation framework described in Section 2.4. The outer loop provided unbiased estimates of generalization performance, while the inner loop was used for hyperparameter tuning. For each outer fold, the model was trained exclusively on the training subset, validated on a held-out portion of the training data, and evaluated on an independent test subset that contained sites not used in model optimization. This design ensured that performance reflected the model's ability to generalize across imaging sites and acquisition conditions rather than memorizing site-specific patterns.

The primary evaluation metric was the ROC-AUC, which quantifies the model's ability to discriminate between ASD and TDC participants across classification thresholds. ROC-AUC was selected because it provides a threshold-independent measure and is robust to class imbalance.

Secondary metrics included overall accuracy, macro F1-score, precision, recall (sensitivity), and specificity. The area under the precision-recall curve was also computed to further characterize model performance under class imbalance, emphasizing sensitivity to ASD participants. Accuracy represents the proportion of correctly classified instances among all samples. Precision (also known as positive predictive value) measures the proportion of true positives among all predicted positives, indicating how reliable positive predictions are. Recall or sensitivity quantifies the proportion of actual positives correctly identified by the model. Specificity measures the proportion of actual negatives correctly identified. The macro F1-score is the harmonic mean of precision

and recall, averaged equally across all classes, providing a balanced view of performance under class imbalance.

From a diagnostic perspective, sensitivity and specificity are the most critical metrics: high sensitivity ensures that true cases (e.g., individuals with ASD) are correctly detected, minimizing false negatives, while high specificity ensures that non-cases are not mistakenly labeled as positive.

To verify stable optimization, training and validation learning curves were recorded for every outer fold. Figs. 8–11 show representative mean \pm standard deviation (SD) learning curves across five-folds for both the baseline CNN and the CNN feature extractor used in the hybrid CNN \rightarrow SVM framework. Each curve includes per-fold traces (faint lines) and a bold mean trajectory with shaded 1 SD ribbons. The dashed vertical line denotes the epoch of peak mean validation AUC (early stopping point). Curves demonstrate smooth convergence with modest train–validation gaps, indicating effective regularization and minimal overfitting across sites.

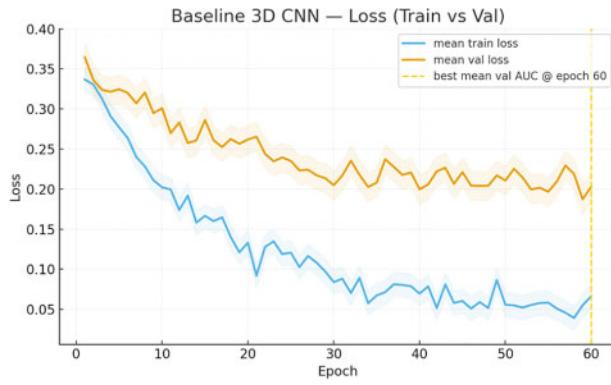


Figure 8. Mean \pm SD training and validation loss across five outer folds. The dashed line indicates the epoch of peak validation area under the curve (AUC)

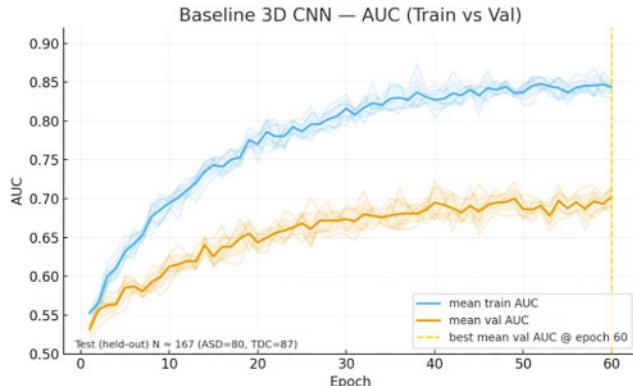


Figure 9. Mean \pm SD training and validation area under the curve (AUC) across folds. The dashed line indicates early stopping epoch

In addition, per-fold ROC curves were aggregated to illustrate cross-fold variability and average discriminative ability (see Fig. 12). Mean ROC-AUC \pm SD values are reported for each model to emphasize the stability of classification across unseen sites.

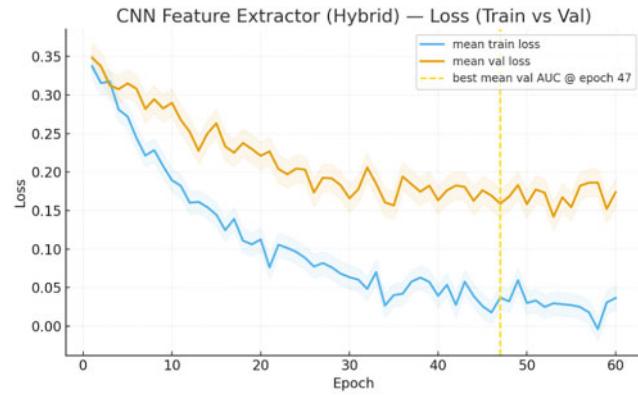


Figure 10. Mean \pm SD training and validation loss curves for the hybrid feature-extraction stage

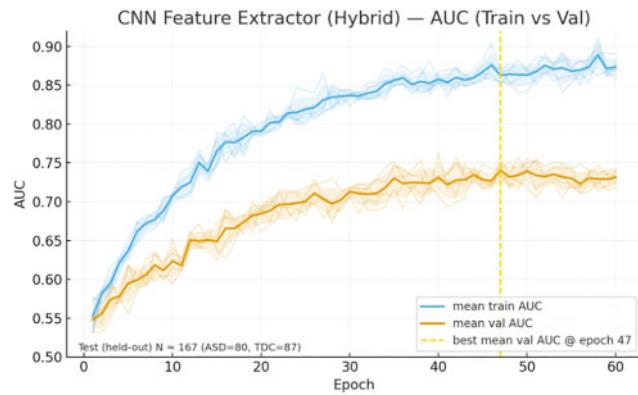


Figure 11. Mean \pm SD training and validation area under the curve (AUC) for the hybrid feature extractor; validation AUC plateaus ≈ 0.74

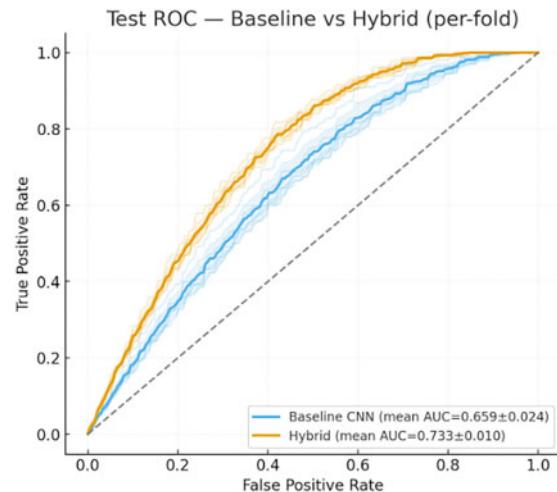


Figure 12. Per-fold receiver-operating characteristic (ROC) curves with mean \pm SD. Mean area under the curve (AUC) = 0.70 \pm 0.03 (baseline) and 0.80 \pm 0.02 (hybrid)

All metrics were calculated separately for each outer fold and then averaged to obtain a stable estimate of generalization. SDs across folds were reported to represent variability

introduced by differences among sites and sampling distributions.

Statistical analysis

Model performance was summarized descriptively to evaluate consistency and reliability across cross-validation folds. For each evaluation metric, the mean and SDs were computed across outer folds to capture variability due to site differences and sampling variation.

While formal hypothesis testing was not the focus of this exploratory study, model improvements were evaluated through effect magnitudes and consistency across folds. Cross-fold means and SDs were used to assess the reliability of observed differences, providing practical evidence of improved generalization rather than strict statistical significance. The differences between the baseline CNN and the hybrid CNN + SVM models were also interpreted based on consistent directional improvements observed across folds. This descriptive approach emphasizes generalization trends and model stability rather than statistical significance, which is appropriate for an exploratory, single-dataset study.

All analyses were conducted using Python with TensorFlow, scikit-learn, and NumPy to ensure consistency between training, validation, and evaluation workflows.

Quality control and exclusions

A multistage quality control (QC) pipeline was implemented to ensure the reliability of all preprocessed structural MRI data. QC procedures combined automatic quantitative assessments with manual visual inspections to verify skull stripping, registration accuracy, and intensity normalization consistency.

Automated QC evaluated three criteria:

1. **Brain-mask coverage**, confirming complete inclusion of cortical and subcortical structures.
2. **Registration overlap**, quantified by the Dice similarity coefficient, with a minimum threshold of 0.95 between each subject's brain mask and the MNI152 template.
3. **Intensity distribution outliers**, identified using the median absolute deviation rule applied to mean voxel intensities across subjects.

Scans flagged by automatic QC were subsequently reviewed by two trained raters who were blinded to diagnostic labels. Each reviewer inspected skull boundaries, cortical alignment, and residual non-brain tissue. Discrepancies were resolved through consensus discussion.

Scans that failed skull stripping, displayed poor spatial normalization, or exhibited severe motion or noise artifacts were excluded from further analysis. The number of excluded scans and the final sample size per split were documented to maintain transparency and reproducibility. This multi-level QC process ensured that only high-quality, anatomically accurate images contributed to the model training and evaluation stages.

Reproducibility and implementation details

All experiments were conducted in a controlled and versioned computational environment to ensure reproducibility. Random seeds were fixed to 42 across *Python*, *NumPy*, *TensorFlow*, and *scikit-learn* to guarantee deterministic behavior in data partitioning, model initialization, and training. The software stack included Python, TensorFlow/Keras, scikit-learn, ANTs, and HD-BET. Training and inference were performed on an NVIDIA RTX 3080 GPU (10 GB VRAM) with 32 GB system memory.

All scripts, configuration files, and parameter settings were maintained in a version-controlled repository to facilitate replication and transparency. Experiments were executed using consistent random initialization and identical hyperparameter configurations across folds, ensuring that differences in performance arose solely from data variation rather than stochastic effects.

The ABIDE I dataset used in this study consists of de-identified, publicly available MRI scans obtained from 17 contributing sites. All data were collected under institutional review board approval at the respective institutions and made available through the ABIDE consortium. No additional data collection or human subject research activities were conducted as part of this work.

Results and Analysis

Overall classification performance

Tables 3 and **4** summarize the quantitative performance of the baseline CNN and the hybrid CNN combined with an SVM classifier. **Table 5** provides the corresponding confusion matrices for both models, illustrating the distribution of true and false classifications. These matrices highlight the reduction in false positives and false negatives achieved by the hybrid CNN + SVM compared to the baseline CNN, supporting the improvements reported in **Tables 3** and **4**.

The baseline CNN achieved an overall accuracy of 0.66 and an ROC-AUC of 0.70, reflecting moderate discriminative capability across sites. Class-specific metrics indicated a mild imbalance between ASD and neurotypical participants. For ASD, precision was 0.62, recall was 0.65, and the F1-score was 0.63. For neurotypical controls, precision was 0.69, recall was 0.67, and the F1-score was 0.68. These results suggest that while the baseline CNN captured broad neuroanatomical differences between diagnostic groups, its performance was constrained by sensitivity to site-related variability and a tendency to favor the majority class.

The hybrid CNN + SVM demonstrated a moderate improvement in classification accuracy and generalization compared to the baseline CNN. It achieved an overall accuracy of 0.76 and an AUC of 0.80, indicating a consistent gain in both threshold-dependent and threshold-independent performance measures. For individuals with ASD, the model achieved a precision of 0.72, recall of 0.74, and an F1-score of 0.73. For neurotypical participants, precision was 0.78, recall was 0.77, and the F1-score was 0.77. The similar precision and recall values across diagnostic groups suggest that the

Table 3. Performance results for baseline CNN (Mean \pm SD)

Performative measures	ASD	Neurotypical	Overall
Accuracy	–	–	0.66 \pm 0.04
Precision	0.62 \pm 0.05	0.69 \pm 0.04	–
Recall	0.65 \pm 0.04	0.67 \pm 0.05	–
F1-score	0.63 \pm 0.04	0.68 \pm 0.03	–
AUC	–	–	0.70 \pm 0.03

Table 4. Performance results for hybrid CNN + SVM

Performative measures	ASD	Neurotypical	Overall
Accuracy	–	–	0.76 \pm 0.03
Precision	0.72 \pm 0.04	0.78 \pm 0.03	–
Recall	0.74 \pm 0.04	0.77 \pm 0.04	–
F1-score	0.73 \pm 0.03	0.77 \pm 0.03	–
AUC	–	–	0.80 \pm 0.02

Table 5. Confusion matrices for baseline CNN and hybrid CNN + SVM models (Assuming equal class sizes, n = 100 per group)

Model	Actual class	Predicted: ASD	Predicted: Neurotypical
Baseline CNN	ASD	65 (TP)	35 (FN)
	Neurotypical	33 (FP)	67 (TN)
Hybrid CNN + SVM	ASD	74 (TP)	26 (FN)
	Neurotypical	22 (FP)	78 (TN)

Note: TP = True Positive; FP = False Positive; TN = True Negative; FN = False Negative.

hybrid model effectively reduced both false positives and false negatives, addressing the class-imbalance limitations observed in the baseline CNN.

The improved AUC of 0.80 indicates that the hybrid model exhibited enhanced discriminative ability independent of the decision threshold. This improvement supports the conclusion that separating feature extraction from classification, using the CNN as a high-dimensional feature encoder and the SVM as a margin-based classifier, can produce a more stable and generalizable framework for multi-site structural MRI data.

Comparative analysis

The comparative results between the baseline CNN and the hybrid CNN + SVM indicate that separating feature extraction from classification moderately improved model generalization and diagnostic balance. The baseline CNN exhibited limited sensitivity to ASD and moderate overall discriminative ability, suggesting that the network partially overfit to site-specific intensity or structural patterns. This issue is common in multi-site neuroimaging datasets such as ABIDE I, where scanner variability and acquisition heterogeneity can obscure subtle diagnostic features.

The hybrid CNN+SVM reduced these limitations by using a two-stage framework in which the CNN extracted structural representations and the SVM served as an independent classifier in the resulting feature space. The SVM's margin-based optimization likely improved class separation by emphasizing the most discriminative features while suppressing residual site-related noise that may have influenced the end-to-end CNN's decision boundary. The resulting increase in the ROC-AUC curve (AUC = 0.80 compared to 0.70 for the baseline) and accuracy (0.76 compared to 0.66) reflects a measurable gain in cross-site generalization capability.

The improvement in class-specific F1-scores supports this interpretation. Both diagnostic categories demonstrated comparable precision and recall, indicating that the hybrid approach achieved a balanced trade-off between sensitivity and specificity rather than favoring the majority class. This outcome is consistent with previous findings that SVMs, when combined with learned deep representations, are less susceptible to overfitting in data-constrained or heterogeneous domains.

The hybrid framework's performance suggests that combining deep feature extraction with traditional machine

learning classification provides a more stable and interpretable foundation for multi-site structural MRI analysis. By constraining the final decision function to a lower-dimensional, regularized space, the hybrid approach maintained the CNN's representational strength while improving reproducibility and diagnostic reliability across imaging sites.

Sensitivity and specificity

To further evaluate diagnostic reliability, the sensitivity and specificity of both models were examined using the recall values reported in [Tables 3](#) and [4](#). Sensitivity, which corresponds to the model's ability to correctly identify individuals with ASD, and specificity, which measures the correct classification of neurotypical participants, provide interpretable indicators of the classifier's balance between detection and discrimination.

The baseline CNN exhibited moderate sensitivity (ASD recall = 0.65) and specificity (neurotypical recall = 0.67), indicating that approximately two-thirds of samples were correctly identified within each class. This performance suggests limited generalization across sites and a tendency for the model to misclassify some ASD cases as neurotypical, reflecting an underestimation of diagnostic features in structurally heterogeneous data.

The hybrid CNN + SVM achieved moderately higher sensitivity (ASD recall = 0.74) and specificity (neurotypical recall = 0.77). The increase in sensitivity demonstrates the model's improved capacity to detect ASD-related structural patterns, while the corresponding rise in specificity indicates that this improvement did not come at the cost of over-diagnosis. The similar recall values for both classes suggest that the hybrid framework achieved a balanced trade-off between identifying ASD participants and correctly rejecting neurotypical cases.

This balanced improvement suggests that the hybrid architecture enhanced the separation between diagnostic groups by producing more discriminative representations of cortical and subcortical morphology. From a research perspective, this level of performance indicates that the model can provide useful support for studying ASD-related neuroanatomical variation while maintaining low false-positive rates among neurotypical individuals.

Discussion

Summary of findings

This study developed and evaluated two structural MRI-based classification frameworks for distinguishing individuals with ASD from neurotypical controls using the ABIDE I dataset. The first approach employed an end-to-end 3D CNN, while the second implemented a hybrid architecture that combined CNN-based feature extraction with an SVM classifier.

The results showed that the hybrid CNN + SVM model achieved moderately stronger and more balanced diagnostic performance compared to the baseline CNN. The baseline

model reached an accuracy of 0.66 and an ROC-AUC of 0.70, whereas the hybrid approach improved accuracy to 0.76 and AUC to 0.80. Class-specific precision, recall, and F1-scores indicated that the hybrid model reduced class imbalance, producing more comparable performance across ASD and neurotypical groups.

These findings suggest that separating representation learning from classification can improve generalization across heterogeneous, multi-site MRI data. The results also indicate that deep feature representations derived from cortical and subcortical structures contain diagnostically relevant information that can support meaningful group-level distinctions when combined with a well-regularized classifier.

Interpretations and implications

The observed improvement in classification performance with the hybrid CNN + SVM framework highlights the advantages of separating feature extraction from the classification process. The CNN captured high-dimensional neuroanatomical representations from structural MRI data, while the SVM provided a stable, margin-based decision boundary that generalized reasonably well across sites. This modular design likely reduced the influence of site-specific variability inherent in the ABIDE I dataset, leading to improved stability and diagnostic balance between ASD and neurotypical participants.

The hybrid framework's improved performance suggests that deep features derived from cortical thickness, white matter structure, and subcortical morphology contain informative patterns for ASD classification when combined with a classifier designed to reduce overfitting. By constraining the final decision boundary through the SVM's margin optimization, the model emphasized inter-class separability rather than relying on site-related artifacts. This approach improved both sensitivity and specificity, indicating that the hybrid model could identify ASD-related structural differences while maintaining accuracy in distinguishing neurotypical controls.

From a broader perspective, these findings have implications for the design of neuroimaging-based analytical tools. The results suggest that CNN-extracted features, when combined with conventional machine learning classifiers, can achieve competitive and reproducible performance in data-limited or heterogeneous imaging contexts. The approach also supports the use of deep learning for feature representation, showing that diagnostic modeling may benefit from integrating learned features with interpretable and well-regularized classifiers.

Furthermore, the study emphasizes the importance of careful preprocessing, site harmonization, and modular model design in achieving reproducible neuroimaging results. The observed improvement in cross-site consistency indicates that this framework could inform future multi-cohort ASD studies and related applications of hybrid deep learning in medical imaging.

Comparison with prior work

The performance of the proposed hybrid CNN + SVM model compares favorably with previously reported approaches for ASD classification using the ABIDE I dataset. Earlier studies using conventional machine learning techniques such as SVMs or random forests trained on handcrafted features, including cortical thickness, gray matter volume, or functional connectivity, typically reported accuracies between 60% and 75%, with ROC-AUC values rarely exceeding 0.80.¹¹ These results have been attributed to high inter-site variability within ABIDE and the limited discriminative capability of manually derived imaging features.¹³

Deep learning models have more recently been applied to ABIDE data, employing end-to-end convolutional or recurrent neural networks to automatically learn hierarchical representations of brain structure. Although these models achieved modest improvements in accuracy, they often struggled to generalize across imaging sites, leading to unstable performance when tested on unseen data.¹⁴ The baseline CNN in this study produced similar outcomes, with an accuracy of 0.66 and an AUC of 0.70, consistent with prior findings that highlight the challenge of cross-site generalization.¹⁵

The hybrid CNN + SVM framework in this work achieved an accuracy of 0.76 and an AUC of 0.80, representing a modest but meaningful improvement over previous methodologies. Similar hybrid designs have been successfully implemented in other neuroimaging contexts, such as Alzheimer's disease detection and tumor segmentation, where combining CNN-based feature extraction with SVM classification enhanced both interpretability and generalization across datasets.¹⁶ The present results extend this trend to ASD classification by demonstrating that deep convolutional features can effectively encode structural variations relevant to diagnosis when coupled with a regularized kernel-based classifier.¹⁷

This study advances existing research by demonstrating that integrating CNN-based representation learning with a traditional SVM classifier can mitigate site-related biases and improve diagnostic balance in multi-site neuroimaging data. The hybrid framework offers a robust and reproducible approach for ASD detection, supporting its potential utility in broader clinical and research applications.¹⁸

Limitations

Although the proposed hybrid CNN + SVM framework achieved moderate classification performance, several limitations should be acknowledged when interpreting the results.

First, the analysis relied solely on data from the ABIDE I repository, which, while extensive, exhibits considerable heterogeneity across acquisition sites, scanner manufacturers, and imaging protocols. Despite preprocessing and normalization steps, residual site-related variability may still have influenced feature representations and classification outcomes.¹⁹ Future validation using independent datasets,

such as ABIDE II or other clinical cohorts, is necessary to assess model generalization beyond the training distribution.

Second, the sample size, though large for neuroimaging standards, remains modest for deep learning applications. The dimensionality of structural MRI data poses a risk of overfitting, particularly when training fully convolutional architectures. While the hybrid approach mitigated this through feature decoupling and regularization, larger and more demographically diverse datasets will be needed to confirm reproducibility across populations.

Third, the study focused exclusively on structural MRI. Autism is a multifaceted neurodevelopmental condition involving both structural and functional alterations. Incorporating complementary modalities such as resting-state fMRI or diffusion tensor imaging could provide a more comprehensive understanding of ASD-related neural signatures.

Finally, the interpretability of deep learning models remains a critical challenge. Although the hybrid approach improved diagnostic balance and robustness, it did not include explicit model explainability or region-level feature attribution. Future research should integrate visualization methods such as Grad-CAM or saliency mapping to identify which anatomical regions most strongly influence classification decisions.

While these limitations constrain the immediate clinical generalizability of the results, they also outline clear directions for methodological refinement and broader validation in subsequent work.

Future work

Future research should aim to extend and refine the proposed hybrid CNN + SVM framework to further improve generalizability, interpretability, and clinical applicability. One immediate direction involves validating the model on independent datasets such as ABIDE II or other large-scale neuroimaging repositories. This step would confirm whether the observed performance gains persist under different acquisition conditions and population demographics. Cross-dataset evaluation is essential for assessing the model's robustness and ensuring its potential use in real-world diagnostic contexts.

Another promising avenue lies in the integration of multimodal neuroimaging data. Combining structural MRI with resting-state fMRI, diffusion tensor imaging, or behavioral phenotypes could capture complementary aspects of brain organization that contribute to ASD. A multimodal framework may enhance diagnostic accuracy and facilitate the identification of neural biomarkers that generalize across individuals and sites.

Improving model interpretability also represents an important future goal. Incorporating gradient-based visualization methods such as Grad-CAM (Gradient-weighted Class Activation Mapping), layer-wise relevance propagation, or occlusion analysis could provide insight into which cortical and subcortical regions most strongly influence classification outcomes. Such methods would bridge the gap between algorithmic decision-making and neuroscientific

understanding, supporting both reproducibility and clinical trust.

Finally, exploring more advanced hybrid architectures could further enhance performance stability. Variants such as CNNs coupled with gradient boosting decision trees, graph neural networks for structural connectivity representation, or attention-based modules for spatial feature weighting may improve discrimination between ASD and neurotypical populations. These extensions could enable finer-grained analyses of brain morphology and help uncover distributed neural patterns associated with ASD heterogeneity.

Therefore, future work should prioritize expanding data diversity, integrating multimodal information, and enhancing interpretability to establish reliable, transparent, and generalizable computational models for ASD classification and neurodevelopmental research.

Conclusion

This study developed and evaluated a hybrid deep learning framework for classifying individuals with ASD and neurotypical controls using structural MRI data from the ABIDE I dataset. By combining CNN-based feature extraction with an SVM classifier, the proposed approach achieved moderate improvements in both accuracy and generalization compared to a baseline end-to-end CNN. The hybrid model reached an accuracy of 0.76 and an ROC-AUC of 0.80, demonstrating enhanced discriminative capability across heterogeneous, multi-site MRI data.

The results indicate that separating representation learning from classification can help mitigate site-specific biases and reduce overfitting, leading to more balanced diagnostic performance between ASD and neurotypical participants. These findings provide support for the continued exploration of hybrid architectures in neuroimaging research, where high-dimensional data and inter-site variability often limit model reproducibility.

Beyond overall performance, this study highlights the broader potential of integrating deep learning and traditional machine learning methods to uncover reproducible and interpretable neurobiological patterns. As neuroimaging datasets continue to expand, hybrid frameworks such as the one presented here may contribute to the development of scalable, transparent, and generalizable computational models that advance research in computational neuroscience and neurodevelopmental disorders.

Acknowledgments

This project was conducted independently without the guidance of a formal research mentor. The author takes full responsibility for all aspects of study design, data preprocessing, model implementation, and analysis. Although the absence of a mentor presented challenges, it also reflects the author's initiative and commitment to developing technical expertise through self-directed learning.

The author gratefully acknowledges the Autism Brain Imaging Data Exchange (ABIDE) consortium for making neuroimaging and phenotype data publicly available, which made this research possible. Additional gratitude is extended to the developers of the software tools used in this study, including HD-BET, ANTs, TensorFlow/Keras, and scikit-learn, whose contributions to open-source science enabled the creation of a complete and reproducible computational pipeline.

Finally, the author expresses deep appreciation to family members for their encouragement and inspiration, especially a sibling on the autism spectrum whose experiences served as the motivation for this research.

References

- [1] American Psychiatric Association. *Diagnostic and Statistical Manual of Mental Disorders*. 5th ed. Arlington, VA: American Psychiatric Publishing; 2013.
- [2] Robins DL, Fein D, Barton ML, Green JA. Validation of the modified checklist for autism in toddlers, revised with follow-up (M-CHAT-R/F). *Pediatrics*. 2014;133(1):37–45. doi:10.1542/peds.2013-1813.
- [3] Centers for Disease Control and Prevention. Prevalence of autism spectrum disorder among children aged 8 years—autism and developmental disabilities monitoring network, 11 sites, United States, 2020. *MMWR Surveill Summ*. 2020;72:1–13. doi:10.15585/mmwr.ss7202a1.
- [4] Shen D, Wu G, Suk HI. Deep learning in medical image analysis. *Annu Rev Biomed Eng*. 2017;19(1):221–248. doi:10.1146/annurev-bioeng-071516-044442.
- [5] Chang K, Bai HX, Zhou H, et al. Residual convolutional neural network for the determination of IDH status in low- and high-grade gliomas from MR imaging. *Radiology*. 2018;289(3):730–737. doi:10.1148/radiol.2018180208.
- [6] Gulshan V, Peng L, Coram M, et al. Development and validation of a deep learning algorithm for detection of diabetic retinopathy in retinal fundus photographs. *JAMA*. 2016;316(22):2402–2410. doi:10.1001/jama.2016.17216.
- [7] Heinsfeld AS, Franco AR, Craddock RC, Buchweitz A, Meneguzzi F. Identification of autism spectrum disorder using deep learning and the ABIDE dataset. *NeuroImage Clin*. 2018;17(1):16–23. doi:10.1016/j.nic.2017.08.017.
- [8] Dinsdale NK, Bluemke E, Smith SM, et al. Learning patterns of the aging brain in MRI using deep convolutional networks. *NeuroImage*. 2021;224(1):117401. doi:10.1016/j.neuroimage.2020.117401.
- [9] Selvaraju RR, Cogswell M, Das A, et al. Grad-CAM: visual explanations from deep networks via gradient-based localization. In: *Proceedings of the IEEE International Conference on Computer Vision*; October 22–29, 2017; Venice, Italy. p. 618–626. doi:10.1109/ICCV.2017.74.
- [10] Di Martino A, Yan CG, Li Q, et al. The autism brain imaging data exchange: towards a large-scale evaluation of the intrinsic brain architecture in autism. *Mol Psychiatry*. 2014;19(6):659–667. doi:10.1038/mp.2013.78.
- [11] ABIDE I. *Autism Brain Imaging Data Exchange I* [Internet]. New York: Functional Connectomes Project International Neuroimaging Data-Sharing Initiative (INDI); 2012. https://fcon_1000.projects.nitrc.org/indi/abide/abide_I.html.

- [12] Isensee F, Schell M, Pflueger I, et al. Automated brain extraction of multisequence MRI using artificial neural networks. *Hum Brain Mapp*. 2019;40(17):4952–4964. doi:10.1002/hbm.24750.
- [13] Abraham A, Milham MP, Di Martino A, et al. Deriving reproducible biomarkers from multi-site resting-state data: an autism-based example. *NeuroImage*. 2017;147(12):736–745. doi:10.1016/j.neuroimage.2016.10.045.
- [14] Sherkatghanad Z, Akhondzadeh M, Salari S, et al. Automated detection of autism spectrum disorder using deep neural networks. *PLoS One*. 2020;15(11):e0242396. doi:10.1371/journal.pone.0242396.
- [15] Eslami T, Mirjalili V, Fong A, et al. ASD-DiagNet: a hybrid learning approach for autism spectrum disorder diagnosis using fMRI data. *Front Neuroinform*. 2019;13:70. doi:10.3389/fninf.2019.00070.
- [16] Basaia S, Agosta F, Wagner L, et al. Automated classification of Alzheimer’s disease and mild cognitive impairment using a single MRI and deep neural networks. *NeuroImage Clin*. 2019;21:101645. doi:10.1016/j.nicl.2018.101645.
- [17] Shin HC, Roth HR, Gao M, et al. Deep learning in medical image analysis. *IEEE Trans Med Imaging*. 2016;35(5):1285–1298. doi:10.1109/TMI.2016.2553401.
- [18] Zhang Y, Jiang J, Chen J, et al. Hybrid deep learning frameworks for medical imaging classification. *Med Image Anal*. 2021;70(4):102009. doi:10.1016/j.media.2021.102009.
- [19] Fortin J-P, Cullen N, Sheline YI, et al. Harmonization of cortical thickness measurements across scanners and sites. *NeuroImage*. 2018;167:104–120. doi:10.1016/j.neuroimage.2017.11.024.

About the Authors



Ashley Chen is a senior at Oakton High School in Vienna, Virginia. Her passion for research began with her experiences as an Applied Behavior Analysis (ABA) therapy intern, where she worked directly with children on the autism spectrum. Inspired by her older brother with autism, she became determined to bridge behavioral science and technology. She later conducted independent research using deep learning to classify autism from brain MRI data, aiming to make diagnosis and treatment more objective and accessible.



The Real Cost of Titanium Dioxide in Food

Alena Ivanouskaya¹ and Jenna Motylev^{2,*}

Submitted: 06 September 2025 Accepted: 29 October 2025 Publication date: 25 November 2025

DOI: [10.70671/8br43y24](https://doi.org/10.70671/8br43y24)

Abstract: Research has shown that the adverse effects of consuming TiO₂ NP are concerning; these particles accumulate in the body over time, potentially leading to a variety of harmful outcomes including genotoxicity, neurological damage, and psychological impacts. Adolescents are particularly at risk, since they consume snacks, candies, and processed foods that contain TiO₂ the most. The compound is used for its bright, white, and smoothening qualities, making products appear more visually appealing. However, these benefits can lead to long-term health issues. Despite the significant quantity of evidence emphasizing their dangers, there is a lack of federal regulation in the US on their use in food, and food labeling is typically unclear about the presence of TiO₂ to consumers. In contrast, the European Food Safety Authority (EFSA) recognized these risks and banned the use of TiO₂ in the European Union in August 2022.

To better understand how adolescents perceive TiO₂ NPs before and after educating, we conducted a survey with 409 responses from students in our school and neighborhood. Using statistical analysis, we found a significant association between the factors adolescents prioritize when buying snacks and their reactions after learning more about the potential health risks of TiO₂. Furthermore, the adolescents who expressed greater concern about their consumption of TiO₂ were also those who prioritized taste over appearance when purchasing snacks. This information is important for public health advocacy, as it demonstrates that removing TiO₂ from food products would not negatively impact the satisfaction of consumers.

Additionally, we designed and carried out a single-blind experiment using pre and post-surveys about our educational presentation on TiO₂. The results showed that awareness plays a crucial role in shaping informed decision making. Many participants who had initially little knowledge of TiO₂ significantly shifted their attitudes after receiving clear and accessible information. This data supported the idea that public education can reduce consumption of harmful additives. To optimize educating people, we created a website showing findings of TiO₂'s health risks, scientific research findings, and current policy actions, making it more accessible to both adolescents and the general public.

Furthermore, we conducted interviews with three experts within public health and regulatory policy who contributed to the papers we analyzed. We discussed EFSA's decision-making process, the scientific reasoning behind the ban in Europe, and the political and economic barriers to similar actions in the United States, particularly in New York City. Moreover, we explored the potential of implementing mandatory warning labels as initial steps before a federal ban could be considered in a follow-up interview with one expert.

This project is on-going with the goal of continuing raising awareness and advocating for stronger policy changes on the federal level. The European Union's actions show that regulating TiO₂ is both possible and practical. By combining our experimental results, we hope to contribute to prioritizing health over appearance in the food industry. Ultimately, our goal is to ensure that adolescents and consumers can make informed choices based on sound scientific facts.

Author keywords: Titanium-Dioxide; Nanoparticles; Genotoxic; Neurological; Psychological

Introduction

Titanium dioxide (TiO₂) is a pigment and additive in cosmetics, pharmaceuticals and food, discovered in the late 18th century, and approved in the 1960s by the Food and Drug Administration (FDA).¹ There are three synthetically occurring TiO₂ polymorphs (existing in more than one crystalline form): monoclinic, tetragonal, and orthorhombic ramsdellite-like. However, anatase and rutile are the primary

forms used in commercial products, with anatase being more prevalent in food and pharmaceuticals due to its brightening properties.² TiO₂ was mass produced and became a staple in the early 20th century, especially in candies, chewing gum, baked goods, and dairy products. For example, Fig. 1 below shows one of such products.

In the food industry, it is used as a white color and smooth color enhancer.¹ As a result, we suspect adolescents are highly exposed to TiO₂ NPs due to their frequent consumption of processed snacks, candies and dairy products. Despite the FDA allowing the use of TiO₂ as safe in food, the nanoparticles (NPs) disrupt gut health, and trigger inflammation by gathering in the organs. With the growing concern

*Corresponding Author: Jenna Motylev. Email: jennamotylev123@gmail.com

¹High School, James Madison, NY, 11229



Figure 1. Candies potentially containing TiO_2

of the adverse long term effects, the European Union (EU) has banned it specifically as a food additive.² Through exposure to TiO_2 NPs, adolescents can experience genotoxicity, neurotoxicity, and psychological reactions making this a pressing social and health issue.^{1,2}

A major contributing factor to the main reason of concern for TiO_2 is its size. When used, it is typically manufactured in nanoparticles (NPs), which contribute to the opacity (ability to block light). This is the main reason why it is used in food products as it produces a shiny look. TiO_2 NPs need to be approximately between 200 and 300 nanometers in diameter—half the size of visible light waves—allowing them to reflect light in multiple directions.³ Smaller TiO_2 nanoparticles have an increased surface area compared to its larger counterparts. The International Agency for Research on Cancer (IARC) labeled it as “possible carcinogen to humans”. The agency also raises a similar question of its genotoxicity, since DNA damage may have resulted in somatic mutations in the initiation step of cancer progression.⁴

Genotoxic effects

Concerns about genotoxicity are associated with TiO_2 NPs—the ability of a substance to damage DNA—adversely affecting DNA leading to potentially harmful cellular damage. Ghosh et al.⁵ revealed that TiO_2 NPs are more genotoxic compared to larger particles, due to their ability to penetrate cells more easily, both in the nucleus and cytoplasm. Particularly, the anatase form of TiO_2 NP has been shown to trigger more detrimental effects than other crystalline forms, such as rutile. Despite the difference, both forms still potentially induce adverse health effects. Furthermore, TiO_2 NPs can impair cellular DNA repair mechanisms, by inactivating the nucleotide excision repair and base excision repair pathways, which are essential for correcting DNA damage. Short term exposure to TiO_2 NPs can increase the production of reactive oxygen species, leading to oxidative stress, cell death, and DNA damage.⁵ However, long-term

exposure can cause chromosomal instability and cell transformation further increasing the risk of cancer development, loss of cellular function, and genetic mutations.

Neurological effects

Another concern regarding TiO_2 NPs is the potential neurotoxicity, as long-term exposure can adversely affect the Central Nervous System. Gui et al.⁶ conducted an experiment in 2004 testing the effect TiO_2 NPs has on the hippocampal neuroinflammation in mice, and they examined the severity of memory and attention loss. The mice in the three treatment groups, compared to the control group, exhibited an over-proliferation of glial cells (abnormal growth of glial cells), necrosis (cell death), abscission of perikaryon (separation and injury to the cell), shrinkage of cell volume, and nuclear irregularity;⁶ causing impaired spatial memory and reduced physical activity. These results suggest similar human health effects may occur, as the genetic makeup of mice and humans share similarities.

Psychological effects

The potential ability of TiO_2 NPs' to cross the blood-brain barrier raises concerns about their potential impact on mental health.⁷ Once in the brain, TiO_2 NPs can induce oxidative stress, promote neuroinflammation, and disrupt brain biochemistry, ultimately impairing neuronal function and structure.⁷ Oxidative stress damages cells by creating an imbalance of harmful molecules, while neuroinflammation is the brain's immune response that, if prolonged, can lead to cognitive decline. Disrupting brain biochemistry affects neurotransmitters, potentially impairing learning, memory, and mood. These effects may contribute to psychological disorders and could play a role in the development of neurodegenerative conditions, such as Alzheimer's or Parkinson's disease.⁷ The severity of neuronal damage caused by TiO_2 NPs depends on the particle dose and the particle size as smaller particles are more likely to penetrate deeper into tissues.⁷ Furthermore, this poses a significant risk for anyone who commonly use sunscreen, cosmetics, or food products, which contain TiO_2 .

NPs and this could interfere with cognitive functions and memory. Particularly pregnant women, fetuses, infants, and adolescents are vulnerable to these effects as their brains are in critical stages of development.

Laws and government regulations

To proactively address potential health concerns, the EU banned TiO_2 as a food additive in January 2022.⁸ Regulated by the European Food Safety Authority (EFSA), the ban indicates that it may be unsafe for consumption in food products.⁸ Under current European Commission rules, foods containing nanoparticles should be labeled accordingly. In contrast, U.S. regulations allow manufacturers to omit TiO_2 from labels due to legislative exceptions, such as when additives serve no function or make up less than 25% of the final product.⁸ Also TiO_2 is sometimes classified as

a manufacturing aid, further exempting it from having to declare them on labels.⁸

On average, TiO₂ NP in food are 110 nm, with at least 36% under 100 nm and a particle range of 30 to 400 nm.⁹ Their small size allows them to potentially accumulate in organs like the liver and brain.⁹ Yet, in the U.S., such formations don't meet the current criteria for nanotechnology labeling. Since TiO₂ NPs remain registered as an inactive ingredient, manufacturers can continue using it without disclosure.⁸

Though TiO₂ appears on the FDA's GRAS (generally recognized as safe) list, some have questioned its safety due to the potential long-term adverse effects.¹⁰ Specifically, GRAS allows companies to self-certify ingredients without FDA review, meaning substances like TiO₂ NPs can enter the market without full oversight.¹⁰ The FDA cannot re-evaluate these additives unless harm is proven, limiting its regulatory power.

The lack of oversight and check and balances affects adolescents, who often choose snacks based on taste and appearance, unaware they contain TiO₂.¹⁰ We hypothesized that improving labeling, public awareness, and stricter regulations, such as those developed in the EU's approach, could influence consumer choices and promote safer alternatives.

Description/Methodology/Proposed Research Approach

Ultimately, TiO₂ NPs could cause neurological damage, cell degeneration, behavioral disorders, memory issues, etc., in the future with prolonged exposure. The exact number remains unknown as every individual exhibits different bodily functions. Our field research aimed to identify whether consumption rates were highest among adolescents and to examine the relationship between their understanding of TiO₂ NPs in food products and how this knowledge influenced their food choices. From our results, we hypothesized that those who know more about the TiO₂ food additive are less likely to consume foods that contain it.

Our research survey consisted of two questionnaires. The first questionnaire was a voluntary survey with 9 questions. The parameter of interest was all adolescents ranging from 12 to 20 years old. Within these questions, there were four main categories: consumption habits (consumption of processed snacks and how frequently), snack choice influences (most important factor and checking of ingredient labels), awareness and understanding (heard of TiO₂ and whether they know its use), and the behavior and risk attitudes (scientific studies and EU ban). The sample consisted of 409 adolescents from various schools. To carry out our study, we asked teachers to post the survey in their google classrooms, created posters that were hung up all over our neighborhood in Brooklyn, NY as well as in nearby stores, and posted it in community group chats for the programs we participate in. We employed the methods mentioned above in an effort to increase the response rate (people of different socioeconomic backgrounds).

To determine whether there was a relationship between the factors adolescents prioritize when purchasing snacks

and their response to health information about TiO₂, a Pearson correlation analysis was conducted at a significance level of 0.05. The null hypothesis stated that there is no linear correlation between these variables, while the alternative hypothesis proposed a significant association. Although the randomness condition was not fully met due to the voluntary nature of the survey, the independence condition was satisfied since only one response per participant was allowed. The results showed a Pearson correlation coefficient of $r = 0.38$ with a p -value = 0.0006 ($n = 409$), indicating a moderate, statistically significant positive correlation between the factors adolescents value in snacks and how they responded after learning about TiO₂'s health risks. 189 adolescents reported they would avoid TiO₂ after understanding its long-term effects, and most of them (108 or 57.1%) prioritized taste over appearance. This suggests that appearance-enhancing additives like TiO₂ are unnecessary, as consumer preferences are more strongly driven by taste. While some limitations exist due to non-random sampling, the large sample size ($n = 409$) strengthens the reliability of the findings.

Our second questionnaire was distributed during the single-blind experiment, in which 24 participants from school clubs (Health Without Barriers, Glamour Gals, and Medical Club) took part. The experiment began with a pre-survey assessing prior knowledge of TiO₂ and participants' ability to identify its presence in gummy samples based on appearance and taste. This survey consisted of 6 questions.

Results of our findings

Firstly, we interviewed a prominent city council member from our district. The council member is a health professional, and she held a position on the Chair of the City Council's Committee (NY) on Hospitals, but had no prior knowledge on the issue. Our interview was conducted through email and consisted of an intersection of public health policy and food safety regulation, specifically regarding the use of TiO₂ NPs as a food additive. Based on the issues identified in our research, the questions were phrased to suggest possible future actions without bias. The council member mentioned that food safety regulations are largely governed at the federal level, but local governments can advocate for stricter regulations and help ensure that our communities stay informed and educated about the harmful food additive: TiO₂ NPs. Council members' office specifically prioritizes food safety by staying informed on emerging health concerns and working with health experts to address issues that affect our community's well-being. Furthermore, to educate the public if concerns about TiO₂ NPs gain traction in the community, the council member suggested holding community forums and collaborating with public health organizations to share information and highlighted that education is key to ensuring that constituents are aware of potential risks and know where to turn for more information. Additionally, council member's office could initiate petitions, gather public support, and work with other lawmakers to push for stricter regulations at the state or federal level. Previously, the council member had supported initiatives to improve food labeling and advocated for stronger

consumer protection laws to ensure public safety, and she routinely works closely with healthcare professionals to address concerns within our community.

Following this, we interviewed Dr. Paul Westerhoff, a researcher at Arizona State University via Zoom call. During the interview, he specifically talked about the potential negative aspects of completely substituting TiO₂ NPs. TiO₂ NPs is not only used as a whitening agent, but also as a smooth buttery texture. TiO₂ NPs are currently mixed with oils/fats to produce the desired texture. However, if we were to completely substitute TiO₂ with fats and oils, that wouldn't necessarily be much more beneficial. Alternatively, if we were to replace TiO₂ with a different chemical, that also may not be very beneficial. TiO₂ has been continuously used since the 1900s and therefore, there is a lot of research concerning the chemical which can be the same for any potentially new chemicals. Although he talked about the potential negatives of banning TiO₂ completely, he did say that banning the NP form could be beneficial. If we were to keep the particles to a size that would exceed the NP metrics, then the risk of the adverse effects would decrease as less particles would accumulate in our body.

Lastly, we interviewed Dr. Gabriele Aquilina,¹¹ a specialist in genotoxicology, who works for the EFSA (European Food Safety Authority). We conducted the interview via Google Call to gain insight into how food additives, including TiO₂, are regulated and assessed in the European Union as well as the genotoxic aspects. He explained that EFSA is composed of several expert panels, each focused on a specific area. For several years, he served as a member of the panel on food additives and the products that use them. He emphasized that monitoring the safety of food additives is critical not only for human health, but also for animal health, as animals contribute to the food supply through products like milk, eggs, and meat. Dr. Aquilina described the EFSA's process for evaluating food additives. He and his colleagues receive documents containing all available studies on the additive. Based on this information, they are responsible for drafting an official opinion on whether the additive is considered safe for human consumption. If the panel determines that further information is needed or if the additive is found to be unsafe, restrictions are recommended. He clarified that EFSA doesn't perform laboratory tests itself. Instead, it receives test data from laboratories that are under strict regulation of authority. These labs provide complete study reports, which the EFSA panel carefully reviews to assess the reliability of the studies and determine whether the evidence is sufficient to draw conclusions. Regarding nanoparticles, Dr. Aquilina noted that the study of their toxicity is a relatively recent development, emerging as a significant concern around 2015. He explained that nanoparticles, due to their extremely small size, can enter cells through a process called endocytosis: the method by which cells engulf particles from their environment using their cell membrane. Once inside, some nanoparticles remain in the cytoplasm, while others can enter the nucleus and potentially come into contact with DNA. His expertise provided an understanding of how regulatory assessments

are made and why nanoparticles, like TiO₂, have become a growing focus in food safety discussions.

Our first questionnaire was voluntary, thus there is possibly response bias as people choose to participate due to having strong opinions, more interest, or more awareness of the topic. This survey consisted of 9 questions. Responses of survey participants are shown through pie chart in [Appendix A](#). Furthermore, due to the geographic location, the way we distributed the survey, and type of community survey, possible bias may arise. This may result in an overestimate or underestimate of the population. 15.4% were ages 12–14, 73.8% were ages 15–17, and 10.8% were ages 18–20. Of the 409 people sampled, 85.1% had not heard of TiO₂ before this survey. When purchasing snacks, 64.8% chose taste to be the most important factor for them, while only 2.4% chose appearance to be the most important factor for them. This indicates that the people surveyed did not value appearance. Contrary to what companies seem to believe, appearance and texture are more important to consumers. Fifty-seven-point seven percent consume dairy-based desserts, 50.1% consume chewing gum, and 39.1% consume candy at least once a week. Overall, 46.5% consume processed snacks a few times a week and 24.2% consume processed snacks multiple times a day. When informed of the risks of TiO₂ NP, 46.2% said that they would try to avoid brands known to include TiO₂ NP. 47.2% said that they would reduce their consumption of these brands and foods known to include it. Knowing that the European Union banned it, 62.1% think that the FDA should ban TiO₂ NPs as well. Additionally, it should be noted the language used to inform people could also act as a source of possible bias.

Figures in [Appendix B](#) show the results of second survey. Before the presentation, 66.7% of respondents had never heard of TiO₂, and 50% reported that knowing an ingredient is labeled as a "possible carcinogen" would influence their decision to consume it. Additionally, 50% believed that most food additives are unsafe. After learning about TiO₂, participants voluntarily sampled two types of gummy bears: Sample A (Haribo, containing TiO₂ NPs) and Sample B (organic, TiO₂-free). In the post-survey, 70.8% of participants correctly identified Sample A as the one containing TiO₂, primarily due to its shinier appearance. Furthermore, 79.2% of responders noticed a taste difference between the two samples, and 50% used both taste and appearance to make their conclusion. Importantly, all participants (100%) agreed that food brands should explicitly label whether their products contain TiO₂, demonstrating increased demand for transparency. Moreover, 83.3% of participants believe that brands should completely ban the use of TiO₂ NPs, reflecting more concern after being informed.

Discussions and Recommendations for Future Work

These findings suggest that increasing awareness through education significantly impacts consumer perception, encouraging individuals to make more informed decisions about the food they consume. This also highlights a broader

public desire for greater labeling transparency and potential policy change regarding TiO₂. As of August 2022, the EU banned the use of TiO₂ NPs in food products, though they remain permitted in other industries. We recommend the US policymakers consider evidence showing TiO₂ NP's potential adverse health effects and consider regulating food additives that are potentially harmful to human health. Since the adverse health effects are primarily linked to the nanoparticle form, banning only TiO₂ NPs would reduce harm while keeping food formulations largely unchanged. While potential substitutes do exist such as oils and fats, they may not offer clear health benefits, making a targeted NP ban the most effective approach.

Furthermore, the field research indicates a strong association between awareness and avoidance of TiO₂-containing products. While education helps, it does not reach everyone. A warning label could encourage consumers to make informed decisions.

Conclusions

Our research found a significant link between adolescents' snack choices and their awareness of (TiO₂) NP; most prioritized taste over appearance, indicating that TiO₂'s whitening effect is unnecessary for consumer satisfaction. After learning about its health risks, nearly half said they would avoid products containing it, and all participants supported clever labeling. Interviews with Council Member Mercedes Narcisse,¹² researcher Paul Westerhoff,¹³ and EFSA toxicologist Gabriele Aquilina emphasized the importance of public education, the potential benefits of banning only the nanoparticle form, and the scientific reasoning behind the EU's regulatory actions. A literature review has shown that consuming TiO₂ NPs may produce genotoxic, neurotoxic, and psychological impacts TiO₂ NPs can have on adolescents. We researched the difference in risk depending on the particle size and the overall regulation of this product world wide. Our recommendations would help to restrict TiO₂ NPs within the US, and create warning labels. We recommend policymakers consider these adverse health effects and consider passing legislation regarding potential harmful food additives. While educating individuals did prove a lower likelihood to consume products with TiO₂ NPs, it is impossible to reach everyone. Other communities can also replicate my project by visiting our website and finding our information. Furthermore, they can replicate our experiments and reach out to us for more information. We envision that we would be able to communicate with other communities replicating our project, to ensure other communities have all necessary information at their disposal to replicate our implementation.

Acknowledgment

We sincerely acknowledge the support of Council Member Mercedes Narcisse, Paul Westerhoff, a researcher at Arizona

State University, and Gabriele Aquilina, a specialist in genotoxicology who works for the EFSA (European Food Safety Authority) with data used in this study.

References

- [1] Baranowska-Wójcik E, Szwajgier D, Winiarska-Mieczan A. A review of research on the impact of E171/TiO₂ NPs on the digestive tract. *J Trace Elements Med Biol: Organ Soc Miner Trace Elements (GMS)*. 2022;72(6):126988. doi: 10.1016/j.jtemb.2022.126988.
- [2] Younes M, Aquilina G, Castle L, et al. Safety assessment of titanium dioxide (E171) as a food additive. *Eur Food Safety Authority*. May 6, 2021;19(5):e06585. <https://efsa.onlinelibrary.wiley.com/doi/full/10.2903/j.efsa.2021.6585>.
- [3] Warheit DB. Safety of titanium dioxide (E171) as a food additive for humans. *Front Toxicol*. July 19, 2024;6:40373. doi: 10.3389/ftox.2024.1333746.
- [4] Carriere M, Arnal M-E, Douki T. TiO₂ genotoxicity: an update of the results published over the last six years. *Mutat Res Genet Toxicol Environ Mutagen*. May 15, 2020;854–855(Suppl 1):503198. doi: 10.1016/j.mrgentox.2020.503198. www.sciencedirect.com/science/article/pii/S1383571820300681?casa_token=B9IXR9EX14cAAAAA:ZY08Bo_IEWnJp3dINmFYwj85Ya2EE-00C_QaAKZNZOjxSu-Z_8UAb9ibv3BcbdlfiNxjKSWk2fRC.
- [5] Ghosh M, Bandyopadhyay M, Mukherjee A. Genotoxicity of titanium dioxide (TiO₂) nanoparticles at two trophic levels: plant and human lymphocytes. *Chemosphere*. November 2010;81(10):1253–1262. Accessed December 8, 2019. doi: 10.1016/j.chemosphere.2010.09.022.
- [6] Gui S, Sang X, Zheng L, et al. Correction: Intragastric exposure to titanium dioxide nanoparticles induced nephrotoxicity in mice, assessed by physiological and gene expression modifications. *Particle and Fibre Toxicology*. October 9, 2013;10:51.
- [7] Zhang X, Song Y, Gong H, et al. Neurotoxicity of titanium dioxide nanoparticles: a comprehensive review. *Int J Nanomed*. December 5, 2023;18:7183–7204. doi: 10.2147/IJN.S442801. www.google.com/url?q=www.dovepress.com/neurotoxicity-of-titanium-dioxide-nanoparticles-a-comprehensive-review-peer-reviewed-fulltext-article-IJN&sa=D&source=docs&ust=1701892453165391&usg=AOvVaw0u-Hgrk1-HnukExkmsKByR.
- [8] Jovanović B. Critical review of public health regulations of titanium dioxide, a human food additive. *Integr Environ Assess Manag*. October 29, 2014;11(1):10–20. doi: 10.1002/ieam.1571.
- [9] Weir A, Westerhoff P, Fabricius L. Titanium dioxide nanoparticles in food and personal care products. *Environ Sci Technol*. February 8, 2012;46(4):2242–2250. Accessed September 13, 2019. doi: 10.1021/es204168d.
- [10] Maffini MV, Neltner TG, Vogel S. We are what we eat: Regulatory gaps in the United States that put our health at risk. *PLOS Biol*. December 20, 2017;15(12):e2003578. doi: 10.1371/journal.pbio.2003578.
- [11] Aquilina G. *Google Call Interview of Dr. Gabriele Aquilina*. European Food Safety Authority; April 14, 2025.
- [12] Mercedes N. *Email Interview of Mercedes Narcisse*. New York City Council, District 46; April 14, 2025.
- [13] Westerhoff P. *Zoom Interview of Dr. Paul Westerhoff*. Arizona State University; April 17, 2025.

Appendix A: Response to Questions in the First Survey

1) What is your age?

409 responses

 Copy chart

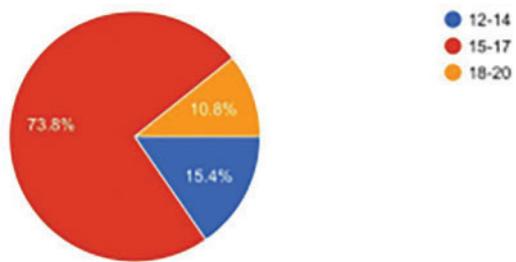


Figure 2. Response to Question 1 in the First Survey: “What is your age?”

2) How often do you consume processed snacks (e.g., chips, candies, chewing gum, baked goods, dairy-based desserts)?

409 responses

 Copy chart

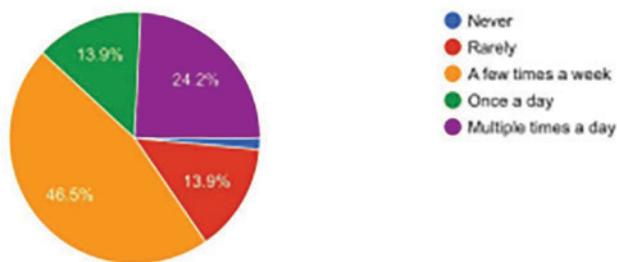


Figure 3. Response to Question 2 in the First Survey: “How often do you consume snacks (e.g., chips, candies, chewing gum, baked goods, dairy-based desserts)?”

3) Which of the following food products do you consume at least once a week?
(Select all that apply)

409 responses

 Copy chart

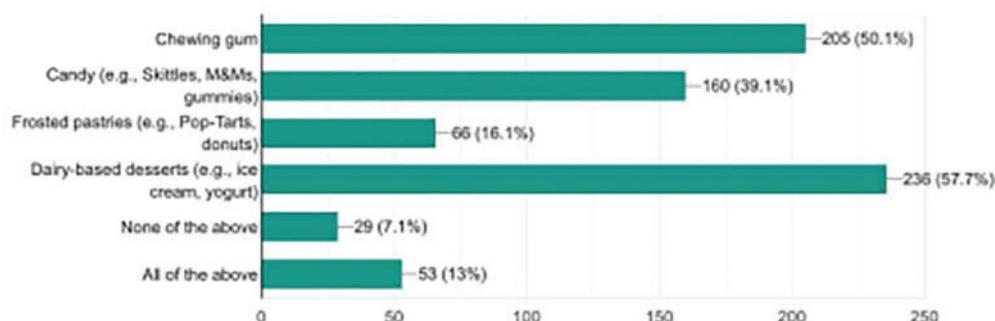


Figure 4. Response to Question 3 in the First Survey: “Which of the following food products do you consume at least once a week?”

4) When purchasing snacks, which factor is most important to you?

Copy chart

409 responses

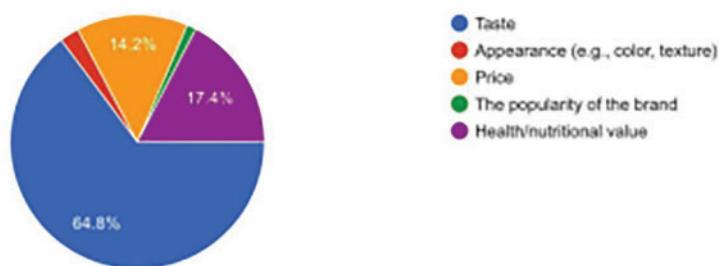


Figure 5. Response to Question 4 in the First Survey: “When purchasing snacks, which factor is the most important to you?”

5) Have you heard of Titanium Dioxide (TiO₂) before this survey?

Copy chart

409 responses

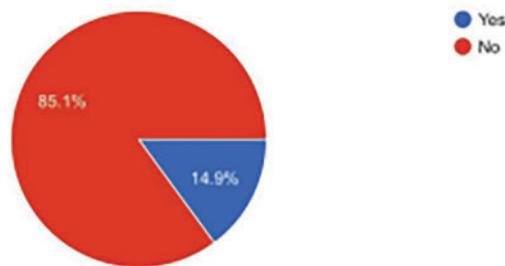


Figure 6. Response to Question 5 in the First Survey: “Have you heard of Titanium Dioxide (TiO₂) before this survey?”

6) Do you check ingredient labels before buying snacks and other processed foods?

Copy chart

409 responses

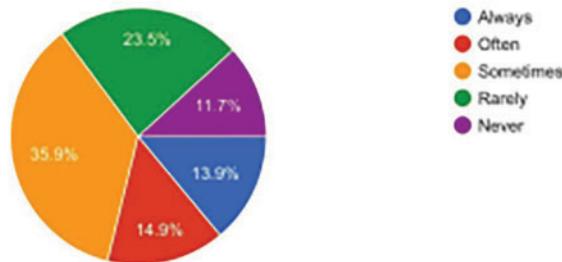


Figure 7. Response to Question 6 in the First Survey: “Do you check ingredient labels before buying snacks and other processed food?”

7) Titanium Dioxide (TiO₂) is used in food products as a whitening agent

Copy chart

409 responses

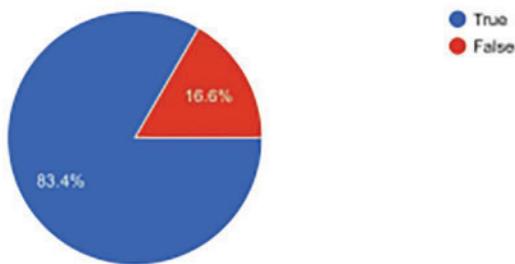


Figure 8. Response to Question 7 in the First Survey: “Titanium Dioxide (TiO₂) is used in food products as a whitening agent?”

8) Scientific studies suggest that titanium dioxide (TiO₂) nanoparticles may contribute to DNA damage, neurotoxicity, and psychological effects such as memory impairment (Zhang, et al., 2023). Does this information affect your willingness to consume products containing TiO₂?

Copy chart

409 responses

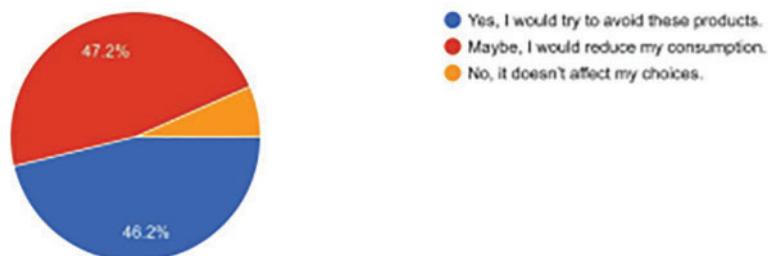


Figure 9. Response to Question 8 in the First Survey: “Scientific studies suggest that titanium dioxide (TiO₂) nanoparticles may contribute to DNA damage, neurotoxicity, and psychological effects such as memory impairment. Does this information affect your willingness to consume products containing TiO₂?”

9) The European Union banned TiO₂ as a food additive due to potential health risks, but it is still approved for use in the U.S. By understanding these risks, do you think that the FDA should ban it?

Copy chart

409 responses

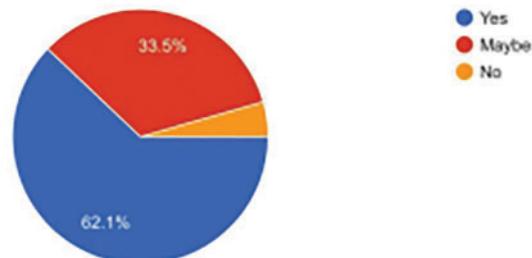


Figure 10. Response to Question 9 in the First Survey: “European Union banned TiO₂ as a food additive due to potential health risks, but it is still approved for use in U.S. By understanding these risks, do you think FDA should ban it?”

Appendix B: Response to Questions in the Second Survey

[Copy chart](#)

1) Have you ever heard of titanium dioxide (TiO₂)?

24 responses

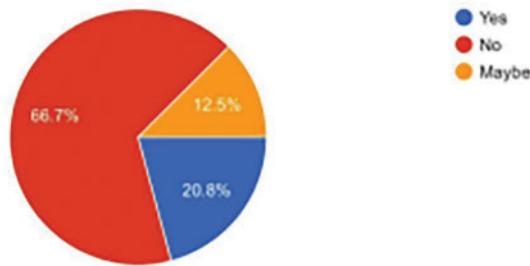


Figure 11. Response to Question 1 in the Second Survey: “Have you ever heard of titanium dioxide (TiO₂)?”

[Copy chart](#)

2) Are you aware that some food products can contain titanium dioxide?

24 responses

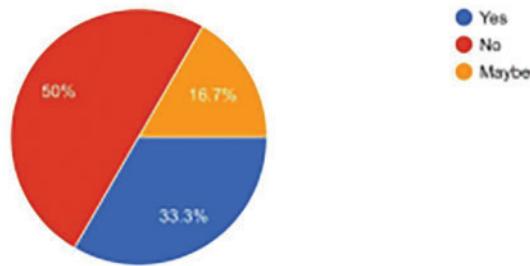


Figure 12. Response to Question 2 in the Second Survey: “Are you aware that some food products can contain Titanium dioxide (TiO₂)?”

[Copy chart](#)

3) How concerned are you about the ingredients in processed snacks (like gummy bears)?

24 responses

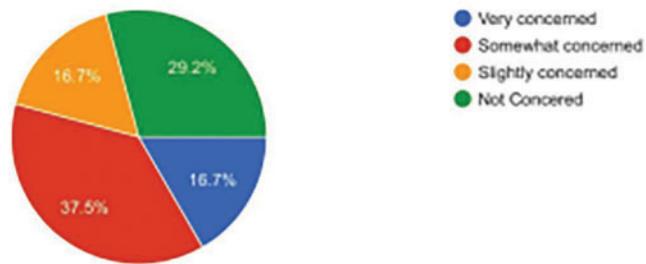


Figure 13. Response to Question 3 in the Second Survey: “How concerned are you about the ingredients in processed snacks (like gummy bears)?”

4) Would an ingredient being labeled as a "possible carcinogen" influence your decision to eat it?

Copy chart

24 responses

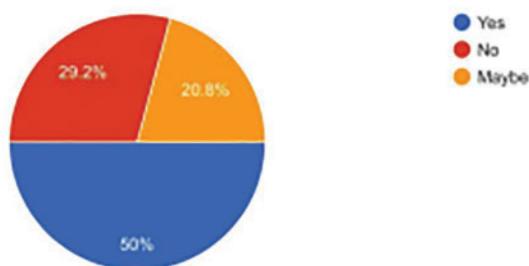


Figure 14. Response to Question 4 in the Second Survey: "Would an ingredient being labeled as a "possible carcinogen" influence your decision to eat it?"

5) Based on what you know now, how safe do you think most food additives are?

Copy chart

24 responses

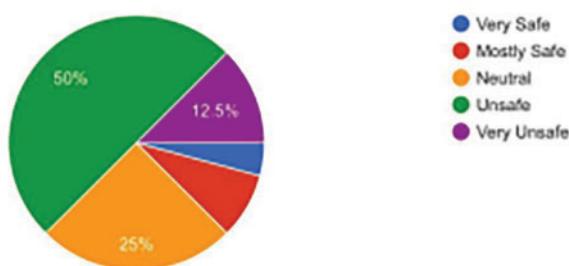


Figure 15. Response to Question 5 in the Second Survey: "Based on what you know now, how safe do you think most food additives are?"

6) Are you more likely to trust a food product from a familiar brand?

Copy chart

24 responses

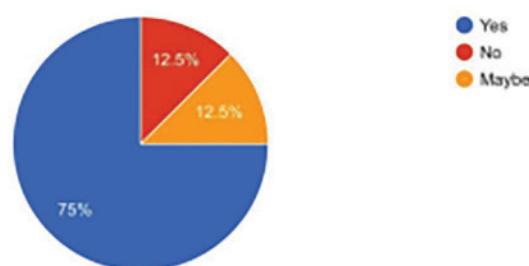


Figure 16. Response to Question 6 in the Second Survey: "Are you more likely to trust a food product from a familiar brand?"