JOURNAL
COMPUTER SCIENCES INSTITUTE

JCSI 39 (2026) 138–145
Received: 21 January 2026
Accepted: 25 April 2026

Statistical analysis of the results of real dice rolls using the object detection model in the context of the Central Limit Theorem

Kacper Gębusia*, Edyta Łukasik

Department of Computer Science, Lublin University of Technology, Nadbystrzycka 36B, 20-618 Lublin, Poland

Abstract

The study conducted an analysis of the convergence of the distribution of sums of real dice-roll results toward the normal distribution in the context of the Central Limit Theorem. The data were obtained from video recordings, and the roll results were identified using the YOLOv8 object detection model. The influence of the original distribution and detection errors on the rate of convergence was examined. A dedicated fitting metric, $S_{n,k}$, based on the least squares method, was employed. The results confirmed convergence to the normal distribution even for small number of samples, achieving a repeatable and stable level of the $S_{n,k}$ metric. The study also demonstrated the robustness of the process to minor data perturbations.

Keywords: Central Limit Theorem; computer vision; analysis of empirical distributions

*Corresponding author

Email address: s95403@pollub.edu.pl (K. Gębusia)

1. Introduction

The Central Limit Theorem (CLT) is regarded as one of the most fundamental theorems in statistics. It underpins numerous methods of data analysis and statistical inference. The theorem states that the sum or mean of a large number of independent, identically distributed random variables with finite mean and variance converges in distribution to the normal distribution.

Theorem 1 (Central Limit Theorem)

Let $X_1, \dots, X_n$ be a sequence of independent and identically distributed random variables with expected value $EX_i = \mu < \infty$ and variance $0 < Var(X_i) = \sigma^2 < \infty$, and let $\bar{X} = \frac{X_1 + \cdots + X_n}{n}$ denote their sample mean. Then, as $n \to \infty$, the standardized variable $\frac{\bar{X} - \mu}{\sigma/\sqrt{n}}$ converges in distribution to a standard normal random variable, that is

$$\lim_{n\to\infty} P(Z_n \le x) = \Phi(x) \quad for\ all\ x \in R$$

where $\Phi(x)$ is the standard normal CDF.

Importantly, the CLT does not require the original distribution of the random variables to be normal. However, several conditions must be satisfied for the theorem to hold:

- the random variables must be independent,
- they must have a finite expected value (mean) and a finite, positive variance,
- the number of variables n must be large (theoretically, $n \to \infty$)[1].

The CLT is a cornerstone of probability theory and mathematical statistics. It provides a mechanism for understanding complex phenomena by approximating their distributions with the well-characterized normal distribution. Yet, it is not merely a theoretical construct; it has broad practical applications in fields such as energy, industry, commerce, social analysis, and demography. For instance, electricity consumption at a given time can be viewed as the sum of consumption across all endpoints in a region. Likewise, numerous small, summable, and typically normally distributed experimental errors contribute to the overall measurement error. The CLT thus formalizes the intuition that the sum of many small, random, and independent effects results in a predictable and stable statistical structure – the normal distribution [2].

It can be said that the CLT lies at the heart of many practical solutions in statistics and data science. However, applied statisticians and data analysts often face unresolved practical questions – most notably, how large n must be for the approximation to the normal distribution to be sufficiently accurate. The theorem itself does not specify the required sample size for reliable inference. Nevertheless, the speed of convergence to normality depends on factors such as skewness and kurtosis of the original distribution [3]. In practice, it is commonly accepted that a sample size of $n \geq 30$ is sufficient in most cases [4].

The history of the CLT dates back to the 18th century. In 1733, Abraham De Moivre derived a normal approximation to the binomial distribution, considered the earliest form of the theorem. Later, in 1810, Pierre-Simon Laplace provided the first proof of the theorem for general discrete distributions. These developments laid the foundation for modern probability theory. Over subsequent decades, the work of mathematicians such as Chebyshev, Markov, Lyapunov, and Lindeberg refined the conditions for convergence and extended the theorem's applicability. The term "Central Limit Theorem" was introduced by Pólya, emphasizing the central role the theorem plays in probability theory [5].

The strength of the CLT lies in the fact that it does not require any specific form for the population distribution. Nevertheless, its abstract formulation may pose comprehension challenges. To facilitate learning, Michael Glencross recommends a hands-on approach involving computer simulations, dice-rolling experiments, and random number sampling [6].

Theoretically, the CLT does not specify a finite number of samples; it only asserts that the number of

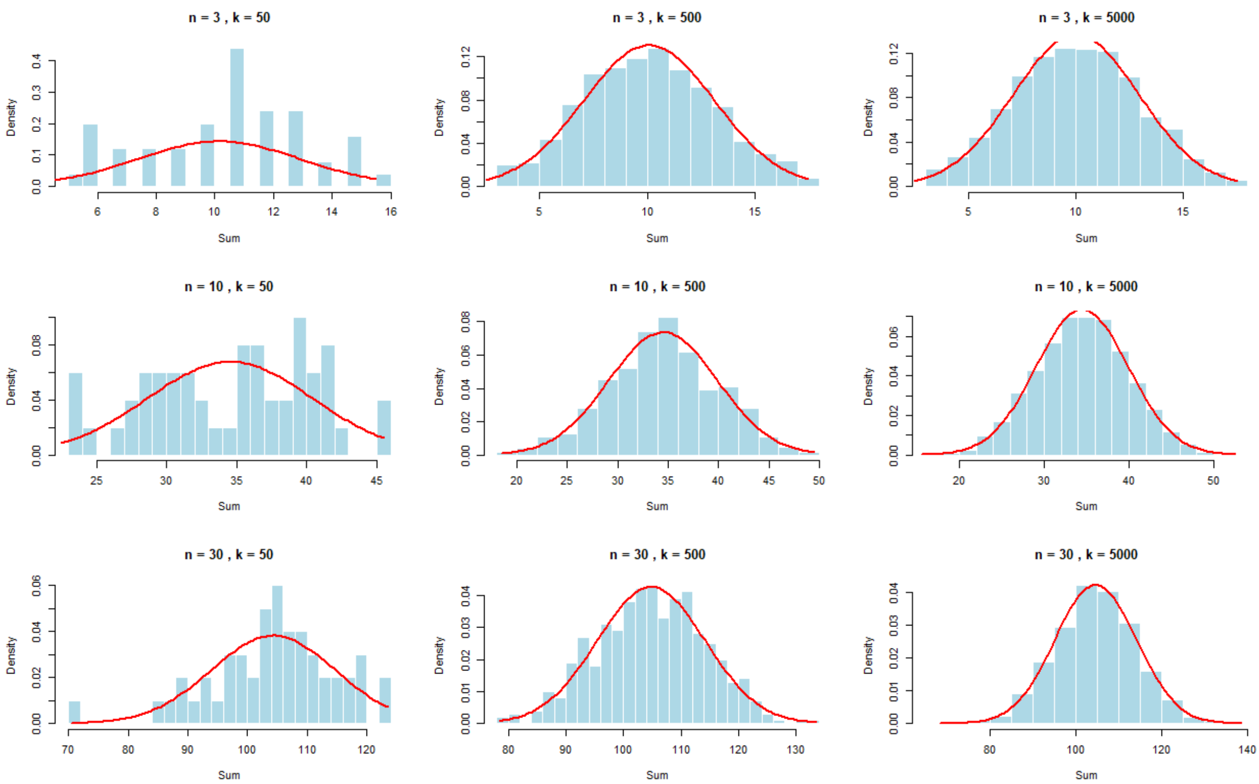

Figure 1: Histograms of sums of random variables generated through simulation

variables should convergent infinity. In practice, this means that the larger the sample size, the more closely the distribution of the sum or mean of the variables will approximate a normal distribution. However, a single sample is insufficient to visually observe the resulting distribution. The extent to which the histogram of sums or means approximates the normal distribution depends not only on the number of sampled variables n, but also on the number of repetitions k in sampling from the population, as illustrated in Figure 1.

A classic visual demonstration of the Central Limit Theorem is the Galton board. It not only illustrates the normal distribution but also serves as an empirical model validating the Central Limit Theorem. Even though the binomial distribution differs from the normal distribution, the distribution of sample means converges to the normal distribution as the number of samples k increases [7].

Online tools and simulators are also available for illustrating the functioning of the Central Limit Theorem. For instance, the process of fitting the frequency distribution of dice roll sums to a normal distribution can be simulated using platforms such as Mathigon. This simulation allows users to observe how, for a fixed number of dice (sample size n), the accuracy of the distribution's fit to the normal distribution improves as the number of trials (number of samples k) increases [8].

The aim of this study is to analyze how the frequency distribution of sums of actual dice roll outcomes, obtained through an object detection model, convergent toward the normal distribution in light of the Central Limit Theorem. The analysis includes an assessment of the rate of this convergence, its accuracy, and the factors that may disrupt it, such as detection errors.

The study is grounded in the formulation of the following hypotheses:

1. H1: The rate of convergence to the normal distribution depends on the original probability distribution of the dice roll results.
2. H2: Increasing the number of observations leads to a more accurate approximation of the normal distribution; however, there is a lower bound on the number of samples required to observe convergence.
3. H3: Detection errors can delay or distort the convergence to the normal distribution.

2. Existing Study

There is an extensive body of literature on the Central Limit Theorem. The literature review focused on analyzing the rate of convergence of dice roll sums to the normal distribution. The article [9] demonstrates that the rate of distribution convergence is highly dependent on the properties of the original distribution. In the case of dependent variables (e.g., martingales), the approximation rate depends on higher-order moments. It has been shown that the higher the values of higher-order moments, the slower the convergence to the normal distribution. Similar conclusions are presented by the author of [10], who determines optimal bounds of convergence depending on the moments of the distribution. Distributions with greater variance or skewness require a greater number of samples to achieve a reliable approximation.

These dependencies are also confirmed by studies on atypical discrete distributions. The article [11] investigates intransitive dice, i.e., sets of dice in which the "wins with" relationship is non-transitive. It proves a Central Limit Theorem describing the distribution of the number of wins of one dice over another as the number of sides increases. The results show that the rate of convergence strongly depends on the value distributions assigned to the individual dice.

Paper [12] analyzes the Central Limit Theorem in the context of distributions built from the sum of digits of numbers in positional numeral systems with base $b \geq 2$. The aim of this article was to prove that the probability measure $\mu(r)$, describing the change in the sum of digits of the number n (a sequence of natural numbers) and the number n after adding a given number r, tends toward a normal distribution. The intended outcome was realized. It was also shown that the variance of the distribution $\mu(r)$ is proportional to the number of blocks, defined as sequences of identical consecutive digits. This means that the more blocks there are, the greater the dispersion of results.

Statisticians typically assume that the sample size should be $n \geq 30$. This number is generally sufficient to meet the assumptions of the Central Limit Theorem. However, sample size should always be tailored to the specific study [13]. The theoretical paper [14] do not provide specific values for minimum sample sizes or the number of samples required to observe convergence to the normal distribution. The author focuses on theoretical convergence in the distribution of sums of random variables to the normal distribution in the limit as the number of variables n tends to infinity. The theorem merely guarantees that, under general conditions, the distribution of means or sums will converge to the normal over time.

There are also studies that describe empirical research. The article [15] analyzes the application of the Central Limit Theorem in casino games, with particular emphasis on variables derived from highly skewed distributions. The study was conducted on samples from simulations. It was shown that the Central Limit Theorem applies even with relatively small sample sizes (~200–500), but larger samples are required for highly asymmetric distributions. At the same time, the study involved a very large number of simulations. This indicates that regardless of the value of n, a large number of samples is necessary to achieve satisfactory convergence to the normal distribution. Similarly, the publication [16] describes a study conducted on simulated data derived from five populations with different distributions. The experiment focused on determining sample sizes n. However, in this case, a much smaller number of samples was used. For specific values of n, the number of samples was only 500. Nevertheless, this was sufficient to observe the operation of the Central Limit Theorem.

The presented studies focused on theoretical considerations or used data generated through simulations. The method of acquiring real-world data may be imperfect, which in turn can limit the applicability of the Central Limit Theorem. The theoretical paper [17] develops the Central Limit Theorem with consideration for data that do not meet ideal assumptions. The authors introduced an index that defines the degree of data distortion. If the level of distortion remains sufficiently low, convergence to the normal distribution in the context of the Central Limit Theorem should still be achieved.

The analyzed studies confirm that the rate of convergence of sample means to the normal distribution depends on the population distribution and that achieving convergence requires not only large sample sizes n but also an adequate number of these samples k. They also show that slightly distorted data do not exclude the possibility of sample means converging to the normal distribution.

3. Material and Methods

The study was designed as an empirical experiment aimed at analyzing the rate of convergence of the frequency distribution of sums from real dice rolls to the theoretical normal distribution. Additionally, the study examines the influence of detection errors and characteristics of the original distribution. To collect the outcomes of dice rolls recorded on video, a script employing an object detection model was utilized.

3.1. Object Detection Model

For the detection of dice roll outcomes, the YOLOv8 object detection model (medium version) was used, trained in the Google Colab environment. The operation of the model is shown in Figure 2.

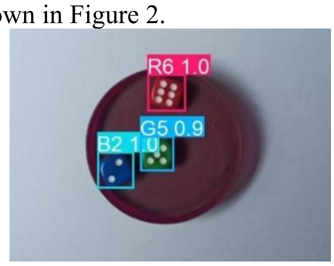

Figure 2: Example detection result obtained during testing of the YOLOv8 model. Bounding boxes indicate detected dice faces and their predicted class labels

Although YOLOv8 demonstrates slightly lower performance compared to some of its predecessors, such as YOLOv5, it offers improved detection capabilities, particularly in challenging conditions [18]. Two object detection models were trained, each characterized by different levels of precision and recall. Model M1 was trained on a correctly labeled dataset. In contrast, model M2 was intentionally degraded by introducing incorrect labels in the training and validation datasets. Specifically, 171 images – representing 28,6% of the total dataset – were deliberately mislabeled. The metric values of the models obtained on the test dataset are presented in Table 1.

Table 1: Precision and recall summary of the trained models

Model	Recall	Precision
M1	98,92%	97,00%
M2	91,93%	83,37%

3.2. Model Dataset

A total of 230 images were taken for the purpose of training, validating and testing the model. The images featured a platform containing three dice in various configurations, as well as an empty field with a designated drop zone for the dice.

The annotation of the images was carried out using the Roboflow platform, which is widely used in computer vision development by providing tools for annotation, training, and model deployment. Roboflow also offers an extensive repository of open-source datasets [19]. A total of 19 object classes were defined for labeling:

- R1 through R6 – representing the respective face values of the red dice,
- G1 through G6 – for the green dice,
- B1 through B6 – for the blue dice,
- X – indicating the area without the platform.

Some of the photos were atypical, e.g., blurry or depicting stacked dice. However, these images were intentionally retained to ensure that the model could generalize to such irregular scenarios. The dataset was split into training, validation and test sets in a 8:1:1 ratio. The data also underwent preprocessing and augmentation. All images were automatically oriented and resized from the original resolution of 3472×3472 pixels to 640×640 pixels. The training dataset was further expanded by applying transformations such as rotation, saturation, blur, and noise. The outcome of these operations was a dataset comprising 598 images, including 23 images allocated for model validation and testing respectively.

3.3. Data Description and Curation

The research data were obtained from recordings of real dice rolls. The rolls were performed in the same manner as those used for training the object detection model. Each recording session lasted between 1 and 50 minutes. The videos were processed using a Python script, which generated a data frame containing the roll number and the outcomes of the individual dice.

Both fair (symmetrical) and loaded dice were used during the experiment. The modification involved weighting the dice on one side. In total, 8,030 rolls were recorded: 6,008 with fair dice and 2,022 with dice exhibiting a biased probability distribution of outcomes. Since each roll involved three dice, the dataset comprises 24,090 individual observations. Due to the fact that the outcomes were recorded automatically, the results of some individual rolls could not be captured. Therefore, the missing values were imputed randomly using the empirical probability distribution estimated from the remaining rolls of the respective dice. The outcomes of the fair-dice rolls recorded by model M1 were divided into three blocks of 2,000 rolls each. The resulting sample sets were used in the study as control samples.

3.4. Experimental Setup

The experiment was conducted on a matte-surfaced worktop to ensure uniform light reflection. The dice rolling area was marked with a black cross inscribed in a circle on a white sheet of paper. A burgundy-colored dice cup with a matching tray was used for rolling. The experiment utilized standard six-sided dice (D6) in red, green and blue. Images and video recordings were captured using a smartphone at various times of the day and under variable lighting conditions, though sufficient for moderate-quality image capture. The smartphone was mounted on a tripod with the lens positioned perpendicularly to the worktop at a distance of 38 cm. The PNG-formatted image capture was handled by a Xiaomi Redmi Note 8 Pro. Photographs used for the training and validating datasets were taken in a 1:1 aspect ratio at a high resolution of 3472x3472. Videos were recorded in vertical orientation in HD resolution (720x1280) at 30 frames per second. Maintaining an identical setup allowed for maximizing the accuracy of result registration while minimizing the size of the model's training dataset.

3.5. Metric

To analyze the experimental results, a dedicated metric was developed to assess the degree of fit between the empirical distribution and the theoretical normal distribution. The adopted index was based on the method of least squares, measuring the square root of sum of squared differences between the observed frequency distribution and the expected normal distribution. This allowed for an unambiguous determination of the minimum number of samples required for the frequency distribution of the sums to closely resemble a normal distribution.

Formally, let:

n – sample size, i.e., the number of dice rolled per trial,
k – numer of samples,
s – possible sum values in a sample, where $s \in \{n, n+1, \dots, 6n\}$,
$\hat{p}_k(s)$ – empirical frequency of sum s among k samples,
T – sum of outcomes in the individual roll:

$$T = \sum_{i=1}^{n} X_i \quad (1)$$

where X_i is the outcome of the i-th dice,
μ – sum mean:

$$\mu = \frac{1}{k}\sum_{i=1}^{k} T_i \quad (2)$$

σ – standard deviation of the sum:

$$\sigma = \sqrt{\frac{1}{k-1}\sum_{i=1}^{k}(T_i - \mu)^2} \quad (3)$$

$p(s;\mu,\sigma)$– theoretical probability of observing sum s under the fitted normal distribution:

$$p(s;\mu,\sigma) = \Phi\left(\frac{s + 0.5 - \mu}{\sigma}\right) - \Phi\left(\frac{s - 0.5 - \mu}{\sigma}\right) \quad (4)$$

where $\Phi(z)$ is the cumulative distribution function of the normal distribution.
Then:

$$S_{n,k} = \sqrt{\sum_{s=n}^{6n}[\hat{p}_k(s) - p(s;\ \mu,\sigma)]^2} \quad (5)$$

The index $S_{n,k}$ is interpreted as the cumulative fitting error between the empirical frequency distribution and theoretical normal distribution, where:

- lower values indicate a better fit, suggesting that the distribution of dice roll sums approximates the normal distribution,
- higher values indicate deviations from normality.

Analyzing the values and stability of $S_{n,k}$ as the number of samples k increases enables the identification of the moment of convergence between the analyzed distributions.

4. Results

4.1. Characteristics of Dice Roll Results

The analysis of the recorded dice roll outcomes is presented in Table 2. It demonstrates that for the experiments conducted using fair dice and the M1 model, the mean value of individual dice outcomes was approximately 3.47, with a standard deviation of about 1.71. It is noteworthy that these values remained stable across all three blocks of 2,000 rolls, indicating that the data volume was sufficient to ensure reproducibility.

Table 2: Descriptive statistics of dice roll outcomes for each dataset (mean, standard deviation, skewness, and kurtosis)

Model Dice type Rolls No	Mean	Standard deviation	Skewness	Kurtosis
M1 Fair 1-2000	3.478	1.722	0.014	1.714
M1 Fair 2001-4000	3.478	1.708	0.041	1.713
M1 Fair 4001-6000	3.470	1.703	0.042	1.723
M1 Loaded 1-2000	3.354	1.726	0.106	1.713
M2 Fair 1-2000	3.830	1.704	-0.172	1.702

The mean for the samples with loaded dice dropped to around 3.35, which is a clear indication that the distribution has shifted towards lower numbers, as a result of modification that involved weighting the dice on the one side. On the other hand, the sample set created by the M2 model showed a mean of about 3.83. This is the furthest difference from the theoretically expected mean value of 3.5. The standard deviation in these two sample sets was also around 1.7, just like in the control ones.

In the case of the first three sample sets, skewness fluctuated around 0 (0.01 – 0.04). It was slightly higher for the asymmetric dice, reaching approximately 0.11. In contrast, the data recorded by the M2 model exhibited negative skewness of roughly -0.17. These values can be observed in the histograms with the superimposed trend line (Figure 3). Kurtosis was stable across all sample sets and amounted to approximately 1.7, which is typical for discrete uniform distributions.

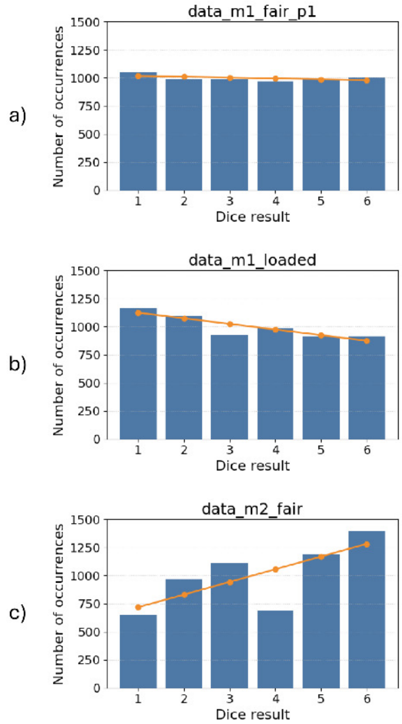

Figure 3: Distributions of individual roll results obtained using fair dice and model M1 (a), loaded dice (b), and model M2 (c)

4.2. Analysis of the $S_{n,k}$ metric

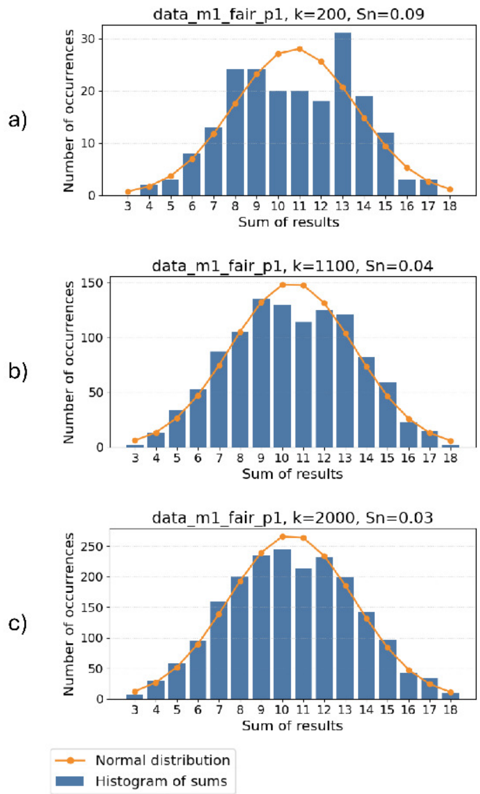

Figure 4: Histograms of the sums of results from fair dice, obtained using model M1 with the superimposed density function of the theoretical normal distribution for $k = 200$ (a), $k = 1100$ (b), $k = 2000$ (c)

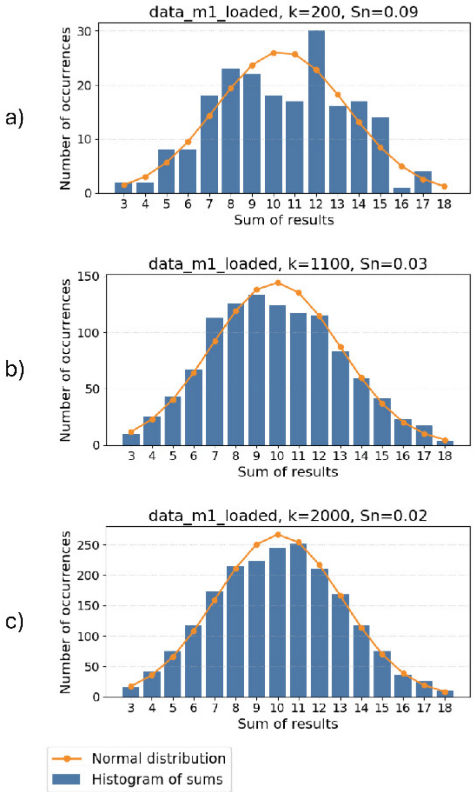

Figure 5: Histograms of the sums of results from loaded dice, obtained using model M1 with the superimposed density function of the theoretical normal distribution for $k = 200$ (a), $k = 1100$ (b), $k = 2000$ (c)

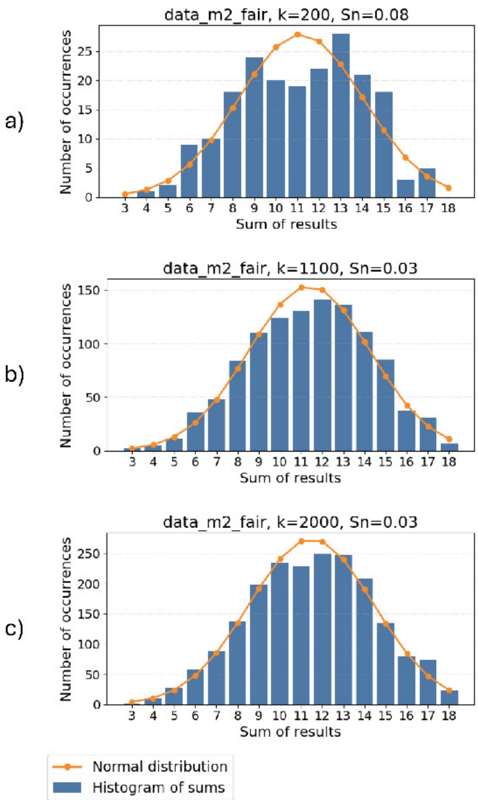

Figure 6: Histograms of the sums of results from fair dice, obtained using model M2 with the superimposed density function of the theoretical normal distribution for $k = 200$ (a), $k = 1100$ (b), $k = 2000$ (c)

The histograms of the sums for the individual sample sets (Figures 4–6) provide a graphical illustration of the degree of conformity to the theoretical normal distribution. Figure 7 presents the plot of the $S_{n,k}$ index for all sample sets as a function of k for $n = 3$. In each case, the decline of the metric occurs at a comparable rate, with the steepest decrease observed within the range $k \in [0, 250]$, followed by a systematic deceleration. For $k \geq 1500$, the $S_{n,k}$ values approach a state of relative stability, and further increases in k yield only marginal benefits. The behavior of the index during this stabilization phase can be examined in greater detail in the plot shown in Figure 8.

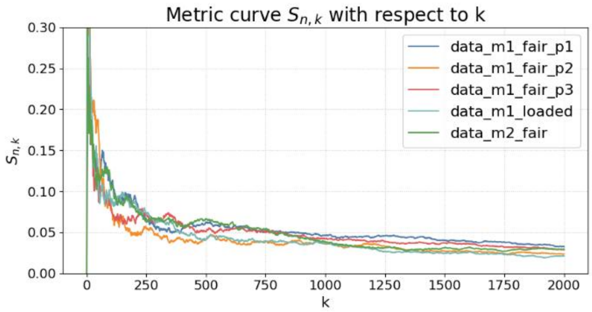

Figure 7: Dependence of the $S_{n,k}$ index on the number of samples k

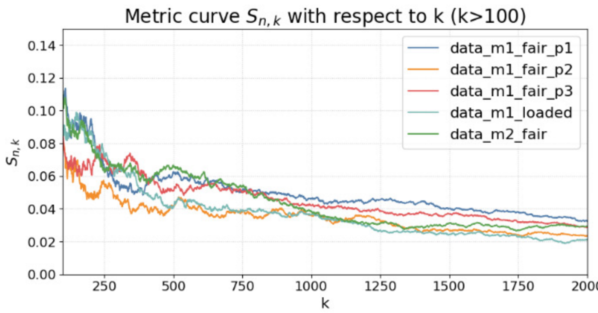

Figure 8: Dependence of the $S_{n,k}$ index on the number of samples k for $k > 100$

No tendency toward faster convergence of the sum frequencies to the normal distribution was observed as the sample set size increased. Differences in the value of $S_{n,k}$ arise from the randomness of the data and occur even among the control sample sets. No single set was identified as exhibiting systematically and substantially lower or higher metric values compared to the others.

5. Discussion

The objective of the study was to empirically verify the rate of convergence of the distribution of sums of real dice-roll outcomes toward the normal distribution in the context of the Central Limit Theorem, using the YOLOv8 object-detection model as a data-extraction tool. The discussion elaborates on the obtained results with respect to the formulated hypotheses and compares them with findings presented in the literature. The limitations of the experiment are also addressed.

5.1. Verification of research hypothesis

H1: The rate of convergence to the normal distribution depends on the original probability distribution of the dice roll results.

The analysis of roll outcomes demonstrated that the sample set obtained using loaded dice differed from the control samples in terms of mean and skewness. At the same time, the analysis of the $S_{n,k}$ indicator did not reveal meaningful differences in the rate of convergence to the normal distribution across trials. The characteristics of the probability distribution of the dice roll results did not translate into a noticeable effect on the degree of conformity to the normal distribution.

H2: Increasing the number of observations leads to a more accurate approximation of the normal distribution; however, there is a lower bound on the number of samples required to observe convergence.

The behavior of the adopted fitting metric indicates that for small values of k the greatest changes in convergence occur. As k increases, the indicator stabilizes in all cases. This hypothesis was confirmed. Indeed, there exists a minimum number of sample that can be considered necessary to observe convergence.

H3: Detection errors can delay or distort the convergence to the normal distribution.

Detection errors were a considerably stronger influence on distortions of the original distribution of roll results than the loaded dice. However, this did not translate into observable deviations in the behavior of the $S_{n,k}$ index in the sample sets employing different detection models. Based on the available data, the hypothesis cannot be confirmed.

5.2. Comparison of the results with the literature

The obtained results diverge from the outcomes one might expect based on the analysis of the existing literature discussed in this article. Publications [13][15][16] emphasize sample sizes on the order of tens, hundreds, or even thousands. In contrast, the present study employed samples of size $n = 3$ and nevertheless succeeded in demonstrating the operation of the Central Limit Theorem.

According to works [9][10][11], the population distribution meaningfully affects the behavior of the theorem. An effect that was not confirmed in the present experiment. Despite differences in population skewness across sample sets, no differences in the rate of convergence were observed.

Conversely, the study [17] shows that, given a sufficiently low level of data contamination, convergence of sample means to the normal distribution is not precluded. The results of this study support that conclusion. Thus, the findings provide empirical evidence that, under the necessary theoretical conditions [1] the Central Limit Theorem holds regardless of the original distribution of the random variables.

5.3. Potential sources of error

The potential limitations of the study arise both from the characteristics of the experiment itself and from the data acquisition method. One such limitation proved to be the insufficient diversity of the analyzed populations. The distributions of the sample sets obtained from the loaded dice and from the M2 model differed too little from the control sample sets. Only relatively small differences in mean values were achieved, along with moderately similar skewness and virtually identical kurtosis. This significantly restricted the ability to draw unequivocal conclusions regarding the influence of the original distribution and detection errors on the rate of convergence in the context of the Central Limit Theorem.

Another important limitation is the low and fixed sample size of $n = 3$. This follows from the nature of the empirical experiment, which assumed that each sample would represent a single physical dice roll. The sum of three dice has a very limited discrete range from 3 to 18, which results in a small number of possible outcomes. Under such conditions, histograms of sums tend to exhibit a more step-like shape when the number of repetitions is relatively small

6. Conclusions and future works

The conducted study enabled an empirical assessment of the convergence of the distribution of sums of real dice-roll results to the theoretical normal distribution in the context of the Central Limit Theorem, using the YOLOv8 object-detection model as a data-extraction tool. The obtained distributions of sums for $n = 3$ and increasing number of samples k exhibited a systematic approach toward the normal distribution, as confirmed by the behavior of the $S_{n,k}$ metric, which stabilized at large values of k. These findings demonstrate that the Central Limit Theorem holds in practice even for small sample size. At the same time, the analysis showed that differences between samples obtained from fair dice, loaded dice, and the M2 model did not translate into a noticeable change in the rate of convergence. In practical terms, this indicates that under the conditions of the experiment, the convergence toward the normal distribution proved robust to minor deviations in the structure of the input data.

The results also highlighted the study's limitations. The most significant of these was the insufficient diversity of the analyzed sample sets, which would be necessary to clear assess the influence of the characteristics of the original distribution on the convergence rate predicted in the literature. Additionally, the analysis was carried out only for a low and fixed number of random variables per sample, $n = 3$.

In light of these findings, future research should focus on expanding the scope of the experiment and increasing the diversity of the analyzed data. A key direction would be the use of an object-detection model that introduces stronger detection errors in favor of specific outcomes. A natural extension involves combining real-world data with computer-generated data, which allows controlled modification of population parameters and enables comparisons of their influence on the rate of convergence. The experiment also relied solely on standard D6 dice, whereas the use of intransitive dice or dice with modified face values could produce populations with more unconventional distributions. Another important direction involves employing different, larger, sample sizes n.

References

[1] X. Zhang, O. L. O. Astivia, E. Kroc, B. D. Zumbo, How to think clearly about the Central Limit Theorem, Psychological Methods 28 (2023) 1427–1445, https://doi.org/10.1037/met0000448.

[2] Q. Cheng, Z. Yang, X. Lu, The Application of Central Limit Theorem, Insight - Physics 5 (2022) 143, https://doi.org/10.18282/ip.v1i1.143.

[3] D. Draper, E. Guo, The practical scope of the central limit theorem, arXiv preprint arXiv:2111.12267 (2021), https://doi.org/10.48550/arXiv.2111.12267.

[4] S. A. Nolan, T. E. Heinzen, Statistics for the Behavioral Sciences, Macmillan, New York, 2014.

[5] W. J. Adams, The Life and Times of the Central Limit Theorem, History of Mathematics, American Mathematical Society, Providence, 2009.

[6] M. J. Glencross, A practical approach to the Central Limit Theorem, In: R. Davidson, J. Swift, (eds) Proceedings of the Second International Conference on Teaching Statistics, International Association for Statistical Education (1986) 91-95.

[7] J. Gong, Galton Board Experiment: Proof of Central Limit Theorem, Science and Technology of Engineering, Chemistry and Environmental Protection 1(3) (2025), https://doi.org/10.61173/j4y96z15.

[8] Mathigon, Probability (Rolling Dice Simulation), https://mathigon.org/step/probability/dice-simulation, [26.11.2025].

[9] L. V. Dung, T. C. Son, On the rate of convergence in the central limit theorem for arrays of random vectors, Statistics & Probability Letters 158 (2020) 108671, https://doi.org/10.1016/j.spl.2019.108671.

[10] I. S. Tyurin, Some optimal bounds in the Central Limit Theorem using zero biasing, Statistics & Probability Letters 82(3) (2012) 514–518, https://doi.org/10.1016/j.spl.2011.11.010.

[11] L. G. Coelho, T. Franco, L. V. Lima, J. P. C. de Paula, J. V. A. Pimenta, G. L. F. Silva, D. Ungaretti, A Central Limit Theorem for Intransitive Dice, arXiv preprint arXiv:2310.17083 (2023), https://doi.org/10.48550/arXiv.2310.17083.

[12] Y. Hosten, É. Janvresse, T. de la Rue, A central limit theorem for the variation of the sum of digits, Annales de l'Institut Henri Poincaré - Probabilités et Statistiques 60(2) (2024) 1125–1149, https://doi.org/10.1214/22-AIHP1346.

[13] V. Kasiulevičius, V. Šapoka, R. Filipavičiūtė, Sample size calculation in epidemiological studies, Gerontologija 7(4) (2006) 225–231.

[14] C. C. Heyde, Multidimensional Central Limit Theorems, In: N. Balakrishnan, T. Colton, B. Everitt, W. Piegorsch, F. Ruggeri, (eds) Wiley StatsRef: Statistics Reference Online (2014), https://doi.org/10.1002/9781118445112.stat02951.

[15] A. Singh, A. F. Lucas, R. J. Dalpatadu, D. J. Murphy, Casino Games and the Central Limit Theorem, UNLV Gaming Research & Review Journal 17(2) (2013) 45–61, https://doi.org/10.9741/2327-8455.1293.

[16] M. A. Barri, A Simulation Showing the Role of Central Limit Theorem in Handling Non-Normal Distributions, American Journal of Educational Research 7(8) (2019) 591–598, https://doi.org/10.12691/education-7-8-8.

[17] B. Berckmoes, G. Molenberghs, Approximate central limit theorems, Journal of Theoretical Probability 31(3) (2018) 1590–1605, https://doi.org/10.1007/s10959-017-0744-6.

[18] E. Casas, L. Ramos, E. Bendek, F. Rivas-Echeverria, YOLOv5 vs. YOLOv8: Performance Benchmarking in Wildfire and Smoke Detection Scenarios, Journal of Image and Graphics 12(2) (2024) 127–136, https://doi.org/10.18178/joig.12.2.127-136.

[19] F. Ciaglia, F. S. Zuppichini, P. Guerrie, M. McQuade, J. Solawetz, Roboflow 100: A Rich, Multi-Domain Object Detection Benchmark, arXiv preprint arXiv:2211.13523 (2022), https://doi.org/10.48550/arXiv.2211.13523.