
Learning Moderately Input-Sensitive Functions: A Case Study in QR Code Decoding

Kazuki Yoda¹ Kazuhiko Kawamoto¹ Hiroshi Kera^{1,2}

Abstract

The hardness of learning a function that attains a target task relates to its input-sensitivity. For example, image classification tasks are input-insensitive as minor corruptions should not affect the classification results, whereas arithmetic and symbolic computation, which has been recently attracting interest, is highly input-sensitive as each input variable connects to the computation results. This study presents the first learning-based Quick Response (QR) code decoding and investigates learning functions of medium sensitivity. Our experiments reveal that Transformers can successfully decode QR codes, even beyond the theoretical error-correction limit, by learning the underlying structure of embedded texts. They generalize from English-rich training data to other languages and even random strings. Moreover, we observe that the Transformer-based QR decoder focuses on data bits while ignoring error-correction bits, suggesting a decoding mechanism distinct from standard QR code readers.

1. Introduction

Over a decade, deep learning has shown remarkable success in learning *input-insensitive* functions that realize the target tasks. Indeed, in most standard tasks, such as image classification, object detection, document summarization, and speech recognition, a slight change in their input (e.g., image perturbation, a few typos) is supposed to have little impact on the output (e.g., classification, document summary). Consequently, for example, data augmentation techniques such as rotation, flipping, and cropping of images enhance the insensitivity (more commonly *robustness*), which leads to substantial performance increases in these tasks.

¹Chiba University ²Zuse Institute Berlin. Correspondence to: Hiroshi Kera <kerah@chiba-u.jp>.

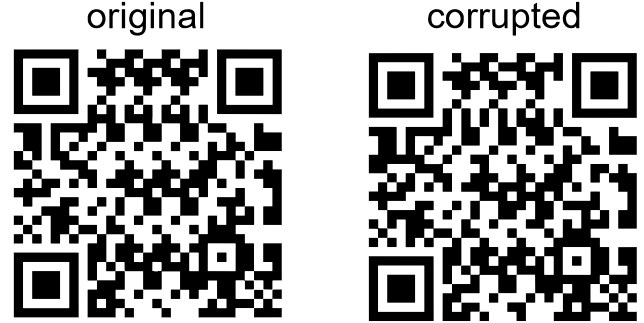


Figure 1. Example of a severely corrupted QR code successfully decoded by the Transformer. “original” denotes the QR code before corruption, and “corrupted” refers to the QR code after applying a 20-bit error. Standard QR code readers fail to decode the corrupted QR code, but the Transformer succeeds. The readers are recommended to try a QR code scan using their smartphone.

Recent studies in arithmetic and symbolic computation further explore the learning of high-sensitivity functions with Transformer models (Vaswani et al., 2017). Examples include integer and modular arithmetic (Power et al., 2022; Shen et al., 2023), symbolic integration (Lample & Charton, 2020), Lyapunov function design (Alfarano et al., 2024), Gröbner basis computation (Kera et al., 2024; 2025), and so on (Wenger et al., 2022; Li et al., 2023a;b; Charton, 2024). In such tasks, changing a single number, coefficient, variable, or operator in the input can immediately change the output, and thus, the target functions to learn are highly sensitive. Although the theoretical difficulty of learning such functions has suggested their hardness (Shalev-Shwartz et al., 2017; Hahn & Rofin, 2024), recent empirical and theoretical studies have shown that high-sensitivity functions can be learned through techniques such as normalization, weight decay, and autoregressive generation (Chiang & Cholak, 2022; Zhou et al., 2024; Kim & Suzuki, 2025). As such, the sensitivity of the target function plays a critical role in determining the difficulty of learning, offering new insights into the capabilities of deep learning models.

In this study, we examine the intermediate case—learning medium-sensitivity functions—through the decoding task of Quick Response codes (QR codes; International Organi-

zation for Standardization (2024)) using a Transformer. QR codes are designed to be tolerant (i.e., insensitive) to corruptions from image capture or physical damage, but decoding them should sensitively capture the change in the plain texts (i.e., embedded strings, such as URLs). Such characteristics place QR code decoding between the low-sensitivity case of classical tasks and the high-sensitivity case of arithmetic and symbolic tasks. QR codes have several error-correction levels and embedding capacities through which we can control their sensitivity. To the best of our knowledge, no prior research has addressed learning-based QR code decoding. There have been many studies for better QR code detection and scanning as a computer vision task (Chou et al., 2015; Kurniawan et al., 2019; Pu et al., 2019; Peng et al., 2020; Shindo et al., 2022; Edula et al., 2023; Zheng et al., 2023; Chen et al., 2024), but these are orthogonal to our context.

Our experiments demonstrate a striking success of Transformer-based decoding under corruptions—the success rate of decoding under corruptions exceeds the theoretical limit, outperforming the standard decoding protocol. Figure 1 shows a case where a standard QR code reader fails to decode the QR code due to severe corruption, whereas the Transformer successfully decodes it. The datasets for training and evaluation consist of QR codes that embed domain names combining random English words and the top 1,000,000 domain names from the Tranco (Pochat et al., 2018). Transformers can retrieve domain names from QR codes even under severe corruptions because of the intrinsic structure of natural language words (e.g., consonants and vowels roughly appear alternately in a word). Surprisingly, the trained Transformer generalizes to non-English words and even random alphabetic strings. Interestingly, the least successful case is the generalization to domain names with irregular words (e.g., “freedoz” instead of “freedom”). In such a case, the Transformer model outputs spell-checked words (i.e., “freedom”), which are incorrect. We also investigated the sensitivity of the functions learned by the Transformer. The Transformer successfully learned functions with four different sensitivity levels, each determined by the error-correction level, demonstrating its ability to decode QR codes. In experiments where the location of errors was controlled, we found that the Transformer learned a function insensitive to the error-correction bits. These findings suggest that the Transformer performs error correction using a mechanism distinct from standard QR code readers.

To summarize, this study addresses a novel QR code decoding task as a showcase of learning medium-sensitivity functions. We empirically obtain the following results through extensive experiments on domain name datasets:

- Transformer-based decoding empirically attains a high success rate with low corruption and maintains moderate success even when the corruption magnitude ex-

ceeds the theoretical limit, which we newly derived for analysis, of error correction.

- Transformers learn natural language structure from English-rich datasets and generalize not only to other languages but even to random alphabetic strings.
- Transformers learn a function sensitive to information essential for decoding while remaining insensitive to redundant information. This tendency suggests that the Transformer performs error correction through a mechanism different from standard QR code readers, which rely explicitly on redundant information.

2. Related Work

Learning High-Sensitivity Functions. Transformers are widely used and excel at learning low-sensitivity functions like image classification (Dosovitskiy et al., 2021), object detection (Zong et al., 2023), and speech recognition (Radford et al., 2023). Recent work shows Transformers can handle high-sensitivity tasks such as arithmetic and symbolic computation (Lample & Charton, 2020; Charton, 2022; Wenger et al., 2022; Li et al., 2023a;b; Charton, 2024; Alfarano et al., 2024; Kera et al., 2024; 2025). For instance, Lample & Charton (2020) demonstrated that Transformers surpass established computational software such as Mathematica and Matlab in accuracy and computational efficiency when solving integrals and ordinary differential equations. Although training on high-sensitivity functions was long considered theoretically difficult (Shalev-Shwartz et al., 2017; Hahn & Rofin, 2024), recent empirical work has shown that techniques such as normalization (Chiang & Cholak, 2022), weight decay (Zhou et al., 2024), and autoregressive generation (Kim & Suzuki, 2025) can successfully surmount these challenges. In previous research, functions with low or high sensitivity have been the focus, leaving medium-sensitivity functions almost unexplored. In this study, we examine Transformer behavior by training models to decode QR codes, thereby probing a function of medium sensitivity.

Deep Learning Approaches to QR Codes. Many studies have explored applying deep learning to QR codes (Chou et al., 2015; Kurniawan et al., 2019; Pu et al., 2019; Peng et al., 2020; Shindo et al., 2022; Edula et al., 2023; Zheng et al., 2023; Chen et al., 2024). For instance, several focus on enhancing detectability by leveraging deep learning models to accurately locate QR codes within images (Chou et al., 2015; Kurniawan et al., 2019; Peng et al., 2020; Chen et al., 2024). On the other hand, other approaches use deep learning models to restore image quality and reduce blurring or low resolution to improve recognition accuracy (Pu et al., 2019; Shindo et al., 2022; Edula et al., 2023; Zheng et al., 2023). QR code reading consists of two principal stages: detection and decoding. Prior works have enhanced reading

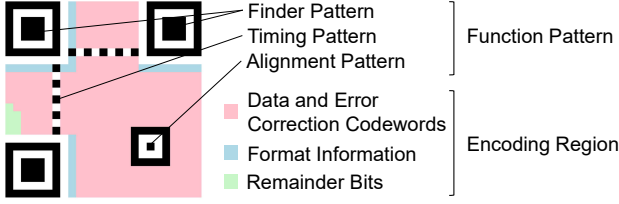


Figure 2. Structure of a v2-QR code. It comprises functional patterns, which support detection, and the encoding region, which embeds the encoded representation of the plain text.

performance by proposing methods to facilitate detection. In contrast, this study aims to analyze the characteristics of the Transformer by training it for decoding.

3. QR Codes

QR codes (International Organization for Standardization, 2024) are two-dimensional matrix codes developed by Denso Wave in 1994 and designed for high-speed information reading. QR codes enable high-speed, accurate transmission of textual information and are used worldwide for various applications, including website access, electronic payments, and airline ticketing.

Figure 2 illustrates the basic structure of a QR code. A QR code consists of black and white units (called *modules*). Some of the units form *function patterns*, which are used at the detection stage by cameras. The rest of the part corresponds to *encoding region*, which encodes the data, along with error correction bits. The encoding region is divided into three components.

Data and Error-Correction Codewords. The plain text is encoded with the Reed–Solomon code and organized into 8-bit units called *codewords*. The data codewords represent the plain text itself, while the error-correction codewords contain the redundant information required for error correction. QR codes support four error correction levels (i.e., L, M, Q, and H in ascending order). As the level increases, the proportion of error-correction codewords to the total number of codewords also increases.

Format Information. It encodes the error-correction level and mask pattern identifier using a BCH code. The resulting bit sequence is duplicated, with one copy placed next to the top-left finder pattern. The other copy is split between the regions next to the bottom-left and top-right finder patterns.

Remainder Bits. These bits are the leftovers resulting from the codeword placement. They contain no actual data or error-correction information and are ignored during decoding.

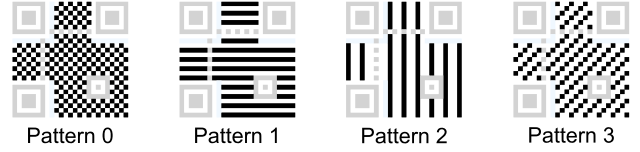


Figure 3. Mask patterns 0 to 3. There are eight mask patterns in total. During masking, black modules within the designated pattern area are inverted. See Appendix A for all mask patterns.

A QR code can have an imbalanced distribution of black and white modules. To mitigate this, a mask pattern is superposed onto the QR code in the final encoding step. There are eight mask patterns (Figure 3), and the most suitable one is selected according to a scoring rule that penalizes poor arrangements, such as long runs of the same color or layouts that hinder detection.

QR codes are defined in 40 versions, ranging from Version 1 to Version 40 by its size. The higher version has more modules, implying higher data capacity. For example, Version 1 has 21 modules per side, Version 2 has 25, and Version 3 has 29. These versions are widely used in product packaging and promotional materials.

In what follows, we write QR codes, for example, of Version 3 and error correction level L as (v3, L)-QR codes.

4. Success Rate Analysis of Error Correction

We will later evaluate the robustness of QR code decoding by Transformers. To this end, we here derive the success rate of error correction in QR codes with n -bit errors. We assume that a bit flip occurs uniformly at random in the encoding region. To the best of our knowledge, this study is the first to derive the theoretical success rate.

Recall that the encoding region consists of the data and error-correction codewords, the format information, and the remainder bits, and the remainder bits are not used in the decoding. Let N denote the total number of bits in the encoding region, and let N_d , N_f and N_r be the numbers of bits in the data and error-correction codewords and the format information, remainder bits, respectively. Accordingly, the numbers of erroneous bits are denoted by p , q , and $n - p - q$, respectively. Then, the successful rate of error correction in encoding region is given by the following.

Theorem 4.1 (Success Rate of Error Correction in the Encoding Region). *Let n be the total number of bit errors in the encoding region. Then, the success rate of error correction—denoted by $P_{\text{success}}(n)$ —is given by:*

$$P_{\text{success}}(n) = \frac{1}{\binom{N}{n}} \sum_{p=0}^n \sum_{q=0}^{n-p} W(p, q) P_d(p) P_f(q), \quad (1)$$

Table 1. Success rate (%) of a model trained on data in which the mask pattern was automatically selected based on the scoring rule, simulating a realistic scenario. The table also shows the proportion (%) of mask patterns in the training dataset. When the scoring rule selects the mask pattern, it tends to favor patterns 1 and 4, resulting in a lower success rate for the other mask patterns.

Mask Pattern	0	1	2	3	4	5	6	7	Average
Success Rate	59.6	92.8	49.1	64.2	90.9	50.7	68.8	54.2	68.3
Proportion	1.3	60.6	1.7	2.7	29.1	1.2	2.4	1.1	-

where P_d and P_f respectively denote the success rate of error correction in the data and error-correction codewords with p -bit errors and that in format information with q -bit errors, and

$$W(p, q) = \binom{N_d}{p} \binom{N_f}{q} \binom{N_r}{n-p-q} \quad (2)$$

is the probability that $(n - p - q)$ -bit errors fall into the remainder bits.

In the following, we provide an overview of the derivation of $P_d(p)$ and $P_f(q)$. All the proofs can be found in Appendix H.

$P_d(p)$ - Success Rate in Data and Error-Correction Codewords. For the data and error-correction codewords, let $M = N_d/8$ denote the total number of codewords, and let M_{ecc} represent the number of error-correction codewords, which depends on the error correction level. According to the properties of Reed–Solomon codes, the maximum number of correctable codewords, denoted by t , is determined by $t = \lfloor M_{ecc}/2 \rfloor$. Under these conditions, the success rate of error correction for the data and error-correction codewords can be determined as follows:

Theorem 4.2 (Success Rate of Error Correction in Data and Error-Correction Codewords). *Let p be the number of erroneous bits within the data and error-correction codewords. Then, the success rate of error correction—denoted by $P_d(p)$ —is given by*

$$P_d(p) = \frac{1}{\binom{N_d}{p}} \sum_{k=\lceil \frac{p}{8} \rceil}^t \binom{M}{k} \sum_{j=0}^k (-1)^j \binom{k}{j} \binom{8(k-j)}{p}. \quad (3)$$

$P_f(q)$ - Success Rate in Format Information. The format information is encoded into 15 bits using a (15, 5) BCH code, and the resulting bit sequence is duplicated so that two identical copies are placed in the QR code. If error correction works in either one of the two, the format information can be correctly read. The (15, 5) BCH code can correct up to 3-bit errors. Under these conditions, the success rate of error correction for format information is given below:

Theorem 4.3 (Success Rate of Error Correction in Format Information). *Let q denote the number of erroneous bits affecting the format information, and let i and j represent the number of bit errors in each of the two respective instances. Then, the success rate of error correction—denoted by P_f —is given by*

$$P_f(q) = \frac{|\{(i, j) \in \mathbb{N}_0^2 \mid i + j = q, \min(i, j) \leq 3\}|}{q + 1}. \quad (4)$$

5. Learning to Decode QR Codes

In this section, we present the evaluation results of the Transformer’s QR code decoding success rate and robustness. We assume that QR code detection has been successfully performed and focus solely on the decoding phase.

5.1. Setup

Task. Transformers are trained for the QR code decoding task. The input is a QR code, and the output is the plain text embedded in the QR code. The QR code is passed to the model in the bit string format.

Dataset. We sampled domain names from Tranco (Pochat et al., 2018), a publicly available ranking of popular domain names. The sampled domain names were encoded into QR codes using the Segno library¹ as a QR code generator. Each QR code was then linearized into a one-dimensional sequence following Figure 4(d). We fixed each dataset’s QR code version, error correction level, and mask pattern. We generated 500,000 samples for the training dataset and 1,000 samples for the evaluation dataset.

Model & Training. To evaluate the general capabilities of the Transformer, we adopted a standard architecture (Vaswani et al., 2017) (six encoder and decoder layers, eight attention heads) and conventional training settings (AdamW optimizer (Loshchilov & Hutter, 2019) with a linearly decaying learning rate starting from 10^{-4}). Following (Lewis et al., 2020), we applied weight sharing between the input embeddings and the output projection layer and employed learnable positional embeddings. The batch size was set to 16, and training was conducted for 10 epochs.

¹<https://github.com/heuer/segno>

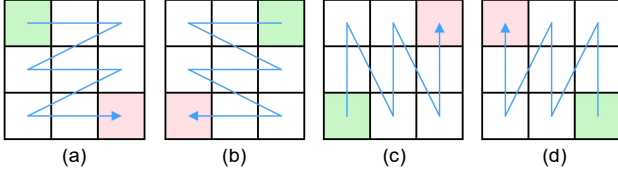


Figure 4. Candidate linearization orders for converting a QR code into a one-dimensional sequence

Table 2. Success rate (%) for four different linearizing orders in (v3, L)-QR codes. The ordering shown in Figure 4(d) outperforms the others, suggesting it is the most suitable for training Transformers.

Order	(a)	(b)	(c)	(d)
Success Rate	93.3	90.2	93.9	95.5

Evaluation. The decoding success rate is defined as the proportion of successfully decoded samples in the evaluation dataset. A decoding is considered successful if the Transformer’s output exactly matches the plain text. To evaluate robustness, we compute the decoding success rate on evaluation datasets composed of artificially corrupted QR codes. We introduce two types of artificial corruption: flip errors and burst errors. Flip errors are introduced by randomly selecting bits within the QR code and inverting their binary values. Burst errors are generated by randomly selecting a 3×3 square region and forcing all bits within that region to be 1. In this experiment, we assume that the detection stage has already been completed. Therefore, corruption is applied only to the encoding region, excluding functional patterns. For evaluation on corrupted data, we compare the Transformer with pyzbar²—a Python wrapper for the widely used open-source ZBar QR code decoder³.

Preliminary Experiments. The aforementioned experimental setup is based on three observations from our preliminary experiments.

1. It was unsuccessful to train Transformers on data where the mask pattern was automatically selected, which reflects a realistic scenario (Table 1). This is because the proportion of mask patterns becomes highly imbalanced. The details can be found in Appendix B.
2. It was easy to train a near-perfect classifier that classifies the mask pattern from input QR codes with almost 100 % accuracy, see Appendix C.
3. Transformer trained on bit strings in the ordering in Figure 4(d) was more successful than in others (Table 2). This is likely because the ordering in Figure 4(d) more

²<https://github.com/NaturalHistoryMuseum/pyzbar>

³<https://github.com/mchehab/zbar>

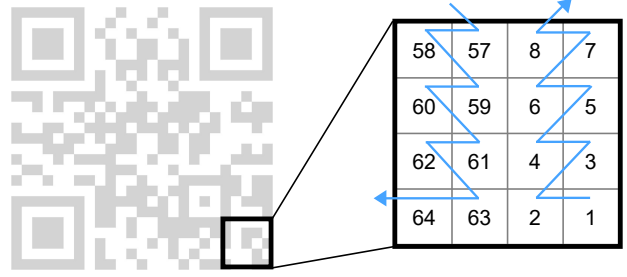


Figure 5. Arrangement order of encoded data in v2-QR Code. The data is placed in a vertical zigzag pattern, moving up and down in two-column bands while avoiding function patterns.

closely matches the bit ordering of QR codes, as illustrated in Figure 5.

From the first two observations, we decided to focus on a fixed mask pattern. From the last observation, we adopted the ordering in Figure 4(d) in our main experiments.

5.2. Decoding Success Rate and Robustness

We evaluated the Transformer’s ability to decode QR codes by decoding success rate, robustness to corruption, and the effect of data augmentation. We also looked at the model’s outputs when decoding both clean and corrupted QR codes.

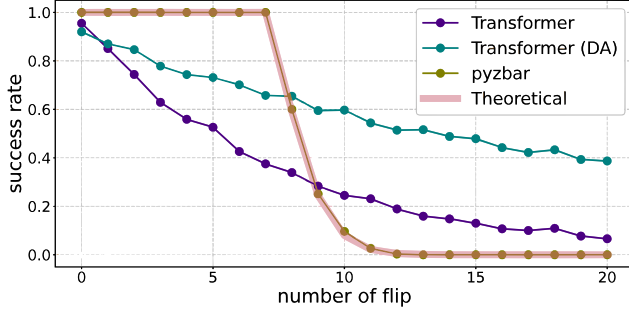
Decoding Success Rate. Table 3 shows the decoding success rate of (v1, L) to (v3, L)-QR codes under each mask pattern. Across all versions, the average success rate exceeded 93 %, indicating high performance. These results suggest that Transformers can learn medium-sensitivity functions.

Decoding Robustness. Figure 6 illustrates the evaluation results for (v3, L)-QR codes (mask pattern 0) with two types of corruption. Figure 6(a) shows that the Transformer achieves a higher success rate than the theoretical success rate of error correction proven in Theorem 4.1 when the number of flip errors exceeds nine. Furthermore, the Transformer surpassed the pyzbar when the number of flip errors exceeded nine or burst errors exceeded four. This result indicates that the Transformer exhibits superior robustness under severe corruption. Additionally, in Appendix E, we present examples of severe corruption where pyzbar fails to decode the QR code but the Transformer succeeds (Figures 10 and 11). Such robustness may be attributed to the model’s ability to learn regularities in word patterns commonly found in domain names, thereby enabling it to infer and reconstruct corrupted information. Note that the plain texts used in the training and evaluation datasets differ.

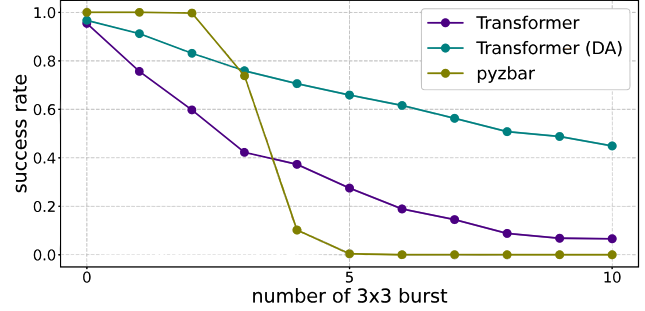
Data Augmentation. We trained the Transformer on (v3, L)-QR codes (mask pattern 0) augmented with flip and burst errors. As shown in Figure 6, it improved significantly

Table 3. Success rate (%) of Transformers for each mask pattern in (v1, L) to (v3, L)-QR codes. Across all versions, the model achieves an average decoding success rate of over 93 %, suggesting that Transformers can learn functions with medium sensitivity. Here, as shown in Section 3, the version determines the number of bits and the structural layout of the QR code; the error correction level represents the degree of redundancy for data recovery; and the mask pattern specifies the type of mask applied after encoding.

Mask Pattern	0	1	2	3	4	5	6	7	Average
Version 1	99.0	97.9	99.2	97.8	98.6	98.5	97.5	97.8	98.3
Version 2	96.0	95.5	96.3	95.0	92.1	95.0	94.5	94.3	94.8
Version 3	95.5	95.5	95.6	93.6	95.1	91.6	93.4	90.8	93.9



(a) Success rate for QR codes corrupted by flip errors



(b) Success rate for QR codes corrupted by burst errors

Figure 6. Success rate on corrupted (v3, L)-QR codes (mask pattern 0) for each method. (a) Success rate under flip noise. (b) Success rate under burst noise. “Transformers (DA)” refers to the model trained with corrupted data augmentation. “Theoretical” refers to the success rate of error correction proved in Theorem 4.1. When the level of corruption is high, the Transformer can decode with a greater success rate than the success rate guaranteed by the QR code’s error correction capability. In addition, applying data augmentation leads to substantial improvements in error resilience.

on corrupted data. These results suggest that training on corrupted inputs can make models more robust.

Examples of Generated Strings. Examples of strings generated during decoding by the Transformer are shown in Table 4. When the generated string matches the original string precisely, the model successfully handled both short and long plain text. In addition, even when the generated strings did not precisely match the original string, they were generally close to the ground truth. For instance, for “dtv2009.gov,” the model outputs “dtv2008.gov,” which differs only by one character. There are also cases where the model makes mistakes because it has learned specific patterns in English words. For example, for the plain text “doggettinc.com,” the Transformer incorrectly generated “doggetting.com.” This was likely because it misinterpreted “inc” as part of the preceding word and segmented it as “getting,” reflecting a misunderstanding influenced by learned linguistic regularities. On the other hand, Table 4(b) shows examples of outputs generated from inputs with 20 random bit flips. Compared to the clean examples in Table 4(a), the outputs are generally less similar to the original strings. However, some errors affect only one character, such as “delicom.global” becoming “dedicom.global.” As shown in Figure 6(a), pyzbar fails to return any output when 20-bit

corruption is applied. This is because it is designed to suppress output entirely once the corruption exceeds a certain threshold to avoid incorrect results. In contrast, the Transformer always produces an output, even if it is incorrect. While this can lead to errors, the outputs often remain close to the original, even with heavy corruption. Additionally, in Appendix E, we present the distribution of similarity between the generated strings and the original plain text under severe corruption (Figure 12).

6. Generalization performance of decoding with Transformer

In this section, we evaluate the generalization performance of the Transformer in QR code decoding on various types of plain text.

6.1. Setup

The training configuration is identical to that in Section 5.1. To evaluate the Transformer’s generalization, we generated eight evaluation sets. Each evaluation sample follows the domain name format $word_1 word_2.tld$, where $word_1$ and $word_2$ are concatenated to form the second-level domain (SLD), and tld denotes the top-level domain (TLD), which

Table 4. Examples of strings generated by the Transformer from clean and corrupted (v3, L)-QR codes (mask pattern 0). For clean inputs (a), the model successfully reconstructs the plain text with a high success rate for both short and long sequences. Even when decoding fails, the generated strings exhibit high similarity to the ground truth. For inputs with 20-bit flips (b), lower-similarity failures are observed. However, despite severe corruption, both successful cases and failures are also present, with high similarity to the ground truth.

(a) Clean QR codes			(b) QR codes with 20-bit flips		
	Ground Truth	Prediction		Ground Truth	Prediction
✓	ellis.ru	ellis.ru	✓	casino-x-dawn.bet	casino-x-dawn.bet
✓	mobile-arsenal.com.ua	mobile-arsenal.com.ua	✓	mobil-isc.de	mobil-isc.de
✓	osakabasketball.jp	osakabasketball.jp	✓	slotsyps.info	slotsyps.info
×	doggettinc.com	doggetting.com	×	delicom.global	dedicom.global
×	elalmanaque.com	elalmanique.com	×	inhand.com	inhcdn.com
×	dtv2009.gov	dtv2008.gov	×	mobile-arsenal.com.ua	mobxle-arsenat.com.mt

is randomly selected from frequent options such as “com,” “org,” or “co.” We constructed eight evaluation sets, each comprising 5,000 (v3, L)-QR codes (mask pattern 0) generated from domain names with varying lexical or structural properties. The eight evaluation datasets are as follows:

- **English, German, Swahili:** words in the respective language.
- **Shuffle:** words created by randomly permuting the characters of English words.
- **Random-alphabet:** random sequences consisting of alphabet characters.
- **Misspelled:** English words in which one character is randomly replaced with an incorrect alphabet.
- **Leetspeak:** English words rewritten using visually similar numerals (e.g., c4t; originally cat).
- **no-TLD:** two English words concatenated without a top-level domain (e.g., pianofox instead of pianofox.com). This structural deviation removes the period and TLD component, setting it apart from the standard domain format used in the other datasets.

6.2. Generalization Performance on Multiple Data Sets

Table 5 presents the evaluation results on eight evaluation sets. The Transformer’s success rate varied depending on the characteristics of the dataset.

Natural Language Datasets. In the English, German, and Swahili datasets, consisting of natural language, the Transformer achieved higher success rates than the others. Since the training data was created from popular domain names in the Tranco, many samples include English words. As a result, the Transformer likely learned common patterns such as frequently used prefixes and suffixes, as well as other frequent patterns like consonant–vowel alternation. The differences in success rates among English, German, and Swahili are likely due to differences in the structure of

words across these languages.

Unstructured Datasets. Compared to the English dataset, the Transformer achieved 3.2 % lower success rates on the Shuffle dataset and 4.1 % lower on the Random-alphabet dataset. This is likely because these datasets lack frequent patterns and regularities typically found in real words.

English-Variant Datasets. Although the Misspelled and Leetspeak datasets retain much of the English word structure, the Transformer achieved lower success rates than the English dataset. Specifically, the success rate dropped by 11.4 % on the Misspelled dataset and by 27.0 % on the Leetspeak dataset. These results suggest that even small changes can cause the Transformer to misread letters and instead predict other characters based on learned patterns. As shown in Appendix D, we identified examples in the Misspelled dataset where the model’s learned regularities of English words led to failures. The extremely low success rate on the Leetspeak dataset may be attributed to the lower frequency of digits than letters in the training dataset.

no-TLD Dataset. The success rate on the no-TLD dataset was notably low at 3.0 %. This can be attributed to the fact that the Transformer was trained exclusively on domain names, which caused it to append a top-level domain during inference in most cases.

The experiment shows that the Transformer learns language rules and frequent patterns during training. As a result, it performs well on datasets made from a natural language with a clear structure. In contrast, it showed a lower success rate on the Shuffle and Random-alphabet datasets, which lack such structure. The Misspelled and Leetspeak datasets kept the basic structure of English words but included small changes. In these cases, the Transformer was often confused by the small changes and mistakenly treated them as clean patterns it had learned before.

Table 5. Success rate (%) across eight datasets. In natural language datasets such as English, German, and Swahili, the Transformer shows a high success rate, indicating that the Transformer effectively leverages linguistic regularities. On the other hand, datasets lacking such structure (e.g., Shuffle, Random-alphabet) or containing minor perturbations (e.g., Misspelled, Leetspeak) result in a lower success rate.

Dataset	English	German	Swahili	Shuffle	Random-alphabet	Misspelled	Leetspeak	no-TLD
Success Rate	99.5	97.1	96.6	95.3	94.4	88.1	72.5	3.0

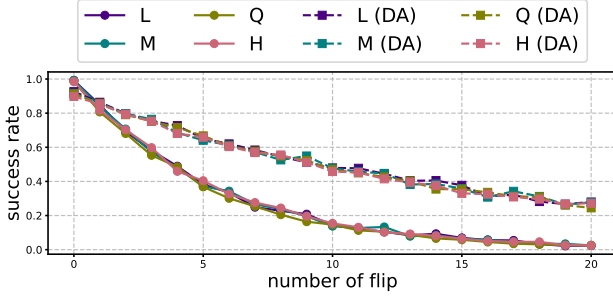


Figure 7. Success rates under (a) flip and (b) burst errors were evaluated for v2-QR codes with all four error correction levels. “(DA)” indicates the result when data augmentation is applied during training. There were almost no differences in success rates between the levels.

7. Sensitivity Analysis of Transformer’s Behavior in QR Code Decoding

In this section, we train Transformers on different error-correction levels (i.e., different input sensitivity) and examine their effect on the trainability of Transformers and the sensitivity after training.

7.1. Setup

Recall that QR codes support four levels of error correction—L, M, Q, and H—in ascending order of strength, with each higher level able to correct more bit errors. Moreover, as the level increases, the proportion of error-correction codewords to the total number of codewords also increases, which leads to lower sensitivity (cf. Table 9). We trained Transformers for the v2-QR code decoding task with each error-correction level. All training and evaluation settings aside from the dataset were identical to those in Section 5.1. The training and test sets were constructed from the same plain texts across different levels so that we could examine the impact of sensitivity differences. Therefore, sensitivity differences appear only in the error-correction codewords.

7.2. Results and Discussion on Sensitivity Behavior

Figure 7 shows the success rates of decoding v2-QR codes across the four error-correction levels. The sensitivity difference did not impact the success rate. To understand this, we plot the success rate decay over the number of flip errors

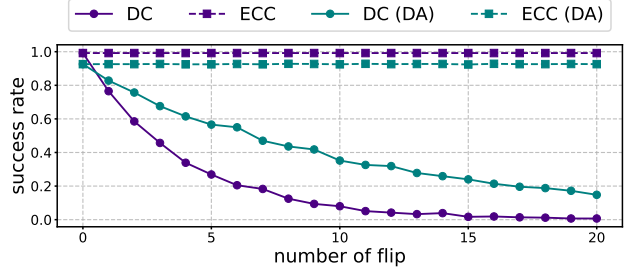


Figure 8. Success rates when bit-flip errors are applied either to the data codewords (DC) or the error-correction codewords (ECC) of (v2, L)-QR codes. “(DA)” indicates the result when data augmentation is applied during training. While errors in the error-correction codewords have little impact, those in the data codewords significantly degrade the success rate.

in Figure 8. It reveals that the bit flips in data codewords decrease the success rate, while those in error-correction codewords do not. This tendency holds even when data augmentation is applied during training. Namely, Transformers do not use error correction bits for QR code decoding.

8. Conclusions

This study tackles the QR code decoding task as an example of a medium-sensitivity function. We demonstrate that Transformers achieve a high success rate under low corruption and maintain moderate decoding success rates even when the corruption exceeds the theoretical error-correction limit, which we newly derived. Moreover, a Transformer trained on datasets rich in English words learns the structure of natural language and generalizes not only to other languages but also to random alphabetic strings. In addition, we show that the Transformer learns a function that is insensitive to error-correction codewords. This finding suggests that the Transformer performs error correction through a mechanism different from that of standard QR code readers.

While our findings provide insights into the Transformer’s ability to decode QR codes as a medium-sensitivity function, this study does not investigate how its internal structure affects decoding performance. In future work, comparing architectures with different positional embeddings and attention mechanisms will be essential to gain a deeper understanding of the Transformer’s behavior in decoding.

Acknowledgments

This research was partially supported by JST PRESTO Grant Number JPMJPR24K4, Mitsubishi Electric Information Technology R&D Center, the Chiba University IAAR Research Support Program and the Program for Forming Japan’s Peak Research Universities (J-PEAKS), and JSPS KAKENHI Grant Number JP23K24914.

References

- Alfarano, A., Charton, F., and Hayat, A. Global lyapunov functions: a long-standing open problem in mathematics, with symbolic transformers. In *Advances in Neural Information Processing Systems*, 2024.
- Charton, F. Linear algebra with transformers. *Transactions on Machine Learning Research*, 2022.
- Charton, F. Learning the greatest common divisor: explaining transformer predictions. In *International Conference on Learning Representations*, 2024.
- Chen, J., Dai, N., Hu, X., and Yuan, Y. A lightweight bar-code detection algorithm based on deep learning. *Applied Sciences*, 2024.
- Chiang, D. and Cholak, P. Overcoming a theoretical limitation of self-attention. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics*, pp. 7654–7664, 2022.
- Chou, T.-H., Ho, C.-S., and Kuo, Y.-F. Qr code detection using convolutional neural networks. *International Conference on Advanced Robotics and Intelligent Systems*, pp. 1–5, 2015.
- Dosovitskiy, A., Beyer, L., Kolesnikov, A., Weissenborn, D., Zhai, X., Unterthiner, T., Dehghani, M., Minderer, M., Heigold, G., Gelly, S., Uszkoreit, J., and Houlsby, N. An image is worth 16x16 words: Transformers for image recognition at scale. In *International Conference on Learning Representations*, 2021.
- Edula, V., Ammisetty, K., Kotha, A., Ravipati, D., Davuluri, A., and P, A. A novel framework for qr code detection and decoding from obscure images using yolo object detection and real-esrgan image enhancement technique. In *International Conference on Computing Communication and Networking Technologies*, pp. 1–6, 2023.
- Hahn, M. and Rofin, M. Why are sensitive functions hard for transformers? In *Proceedings of the Annual Meeting of the Association for Computational Linguistics*, pp. 14973–15008, 2024.
- International Organization for Standardization. Information technology — automatic identification and data capture techniques — qr code bar code symbology specification, 2024.
- Kera, H., Ishihara, Y., Kambe, Y., Vaccon, T., and Yokoyama, K. Learning to compute gröbner bases. In *Advances in Neural Information Processing Systems*, 2024.
- Kera, H., Pelleriti, N., Ishihara, Y., Zimmer, M., and Pokutta, S. Computational algebra with attention: Transformer oracles for border basis algorithms. *arXiv:2505.23696*, 2025.
- Kim, J. and Suzuki, T. Transformers provably solve parity efficiently with chain of thought. In *International Conference on Learning Representations*, 2025.
- Kurniawan, W. C., Okumura, H., Muladi, and Handayani, A. N. An improvement on qr code limit angle detection using convolution neural network. In *International Conference on Electrical, Electronics and Information Engineering*, pp. 234–238, 2019.
- Lample, G. and Charton, F. Deep learning for symbolic mathematics. In *International Conference on Learning Representations*, 2020.
- Levenshtein, V. I. Binary codes capable of correcting deletions, insertions and reversals. *Soviet Physics Doklady*, pp. 707–710, 1966.
- Lewis, M., Liu, Y., Goyal, N., Ghazvininejad, M., Mohamed, A., Levy, O., Stoyanov, V., and Zettlemoyer, L. BART: denoising sequence-to-sequence pre-training for natural language generation, translation, and comprehension. In *Proceedings of the Annual Meeting of the Association for Computational Linguistics*, pp. 7871–7880, 2020.
- Li, C., Wenger, E., Allen-Zhu, Z., Charton, F., and Lauter, K. Salsa verde: a machine learning attack on learning with errors with sparse small secrets. In *Advances in Neural Information Processing Systems*, 2023a.
- Li, C. Y., Sotáková, J., Wenger, E., Malhou, M., Garcelon, E., Charton, F., and Lauter, K. Salsapicante: A machine learning attack on lwe with binary secrets. In *Proceedings of the ACM SIGSAC Conference on Computer and Communications Security*, pp. 2606–2620, 2023b.
- Loshchilov, I. and Hutter, F. Decoupled weight decay regularization. In *International Conference on Learning Representations*, 2019.
- Peng, J., Yuan, S., and Yuan, X. Qr code detection with faster-rcnn based on fpn. In *Artificial Intelligence and Security*, pp. 434–443, 2020.

- Pochat, V. L., Goethem, T. V., Tajalizadehkhoob, S., Korczyński, M., and Joosen, W. Tranco: A research-oriented top sites ranking hardened against manipulation. *arXiv preprint arXiv:1806.01156*, 2018.
- Power, A., Burda, Y., Edwards, H., Babuschkin, I., and Misra, V. Grokking: Generalization beyond overfitting on small algorithmic datasets. *arXiv preprint arXiv:2201.02177*, 2022.
- Pu, H., Fan, M., Yang, J., and Lian, J. Quick response barcode deblurring via doubly convolutional neural network. *Multimedia Tools and Applications*, pp. 897–912, 2019.
- Radford, A., Kim, J. W., Xu, T., Brockman, G., McLeavey, C., and Sutskever, I. Robust speech recognition via large-scale weak supervision. In *International Conference on Machine Learning*, pp. 28492–28518, 2023.
- Shalev-Shwartz, S., Shamir, O., and Shammah, S. Failures of gradient-based deep learning. In *International Conference on Machine Learning*, pp. 3067–3075, 2017.
- Shen, R., Bubeck, S., Eldan, R., Lee, Y. T., Li, Y., and Zhang, Y. Positional description matters for transformers arithmetic. *arXiv:2311.14737*, 2023.
- Shindo, T., Watanabe, T., Yano, R., Arimoto, M., Takahashi, M., and Watanabe, H. Super resolution for qr code images. In *Global Conference on Consumer Electronics*, pp. 274–277, 2022.
- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, L., and Polosukhin, I. Attention is all you need. In *Advances in Neural Information Processing Systems*, pp. 5998–6008, 2017.
- Wenger, E., Chen, M., Charton, F., and Lauter, K. E. SALSA: attacking lattice cryptography with transformers. In *Advances in Neural Information Processing Systems*, 2022.
- Zheng, J., Zhao, R., Lin, Z., Liu, S., Zhu, R., Zhang, Z., Fu, Y., and Lu, J. Ehfp-gan: Edge-enhanced hierarchical feature pyramid network for damaged qr code reconstruction. *Mathematics*, 2023.
- Zhou, H., Bradley, A., Littwin, E., Razin, N., Saremi, O., Susskind, J. M., Bengio, S., and Nakkiran, P. What algorithms can transformers learn? A study in length generalization. In *International Conference on Learning Representations*, 2024.
- Zong, Z., Song, G., and Liu, Y. Detrs with collaborative hybrid assignments training. In *International Conference on Computer Vision*, pp. 6725–6735, 2023.

A. Mask Pattern

QR Code has eight different mask patterns ([International Organization for Standardization, 2024](#)). Figure 9 shows a list of mask patterns.

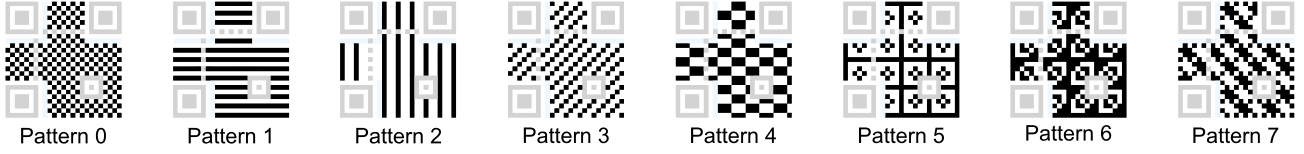


Figure 9. List of mask patterns.

B. Transformer Performance on Data with Mask Patterns Automatically Selected by the Scoring Rule

In QR codes, mask patterns are ordinarily selected automatically by the scoring criteria ([International Organization for Standardization, 2024](#)). In this section, we construct the dataset by encoding v3-QR codes from internet-domain inputs with mask patterns chosen automatically by the scoring rules and subsequently train the Transformer. For evaluation, we prepare 1,000 samples for each of the eight mask patterns, yielding 8,000 samples. The experimental configuration follows that of Section 5.1. As presented in Table 6, the Transformer attains an average decoding success rate of 68.3 % when mask patterns are heterogeneous. Although only patterns 1 and 4 achieve decoding accuracies above 90 %, the remaining patterns fall below the 3 % threshold. We attribute this disparity to the biased proportion of mask patterns within the training dataset. Table 6 depicts this proportion, revealing that most of the 500,000 training samples correspond to patterns 1 and 4. This pronounced imbalance results in an insufficient representation of the other patterns, impairing generalization and reducing inference performance.

Table 6. Success rate (%) of a model trained on data in which the mask pattern was automatically selected based on the scoring rule, simulating a realistic scenario. The table also shows the proportion (%) of mask patterns in the training dataset. When the scoring rule selects the mask pattern, it tends to favor patterns 1 and 4, resulting in a lower success rate for the other mask patterns. (Repeated From Table 1)

Mask Pattern	0	1	2	3	4	5	6	7	Average
Success Rate	59.6	92.8	49.1	64.2	90.9	50.7	68.8	54.2	68.3
Proportion	1.3	60.6	1.7	2.7	29.1	1.2	2.4	1.1	-

C. Mask Pattern Classification via Transformer

To ascertain whether a Transformer can discriminate among mask patterns, we trained and evaluated the model using a dataset encompassing all eight mask patterns. For training, we curated 800,000 popular internet domains from the Tranco list and synthesized the dataset such that each mask pattern was represented by 100,000 samples. For evaluation, we collected 8,000 samples from internet domains distinct from the training set, ensuring each mask pattern comprised 1,000 samples. The architecture and training protocol were identical to those described in Section 5.1. We set the batch size to eight and trained the model for one epoch. The evaluation revealed that the Transformer achieved a classification accuracy of 100 %. This result confirms that mask patterns can be classified with perfect accuracy; consequently, in our proposed method, we assume known mask patterns and train separate models for each pattern.

D. Generated Strings for Misspelled Dataset

Table 7 shows examples of strings generated by the Transformer on the Misspelled dataset in the experiments of Section 6. From the successful cases, it is clear that the Transformer can handle spelling mistakes. However, the failure examples reveal instances where its learned knowledge of English word patterns actually leads to incorrect outputs. For example, the plain text “domajinprotocjl.me” is composed of misspellings of the two English words “domain” and “protocol,” but the

Transformer generated “domainprotocol.me.” This appears to result from misreading “domaij” as “domain” and “protocjl” as “protocol.”

Table 7. Examples of strings generated by the Transformer on the Misspelled dataset in the experiments of Section 6.

	Ground Truth	Prediction	Component Words
✓	drwafcess.io	drwafcess.io	dry + access
✓	eglyiefd.jp	eglyiefd.jp	egg + yield
✓	hopeynww.co	hopeynww.co	honey + new
✓	joqpct.jp	joqpct.jp	job + act
✓	yelloqnobe.io	yelloqnobe.io	yellow + note
×	acrfreedomz.io	acrfreedom <i>m</i> .io	act + freedom
×	domaijprotocjl.me	domainprotocol.me	domain + protocol
×	fameupstre <i>r</i> m.site	fameupstream <i>a</i> .site	flame + upstream
×	gho <i>c</i> tw <i>n</i> reless.fr	ghostw <i>n</i> reless.fr	ghost + wireless
×	ordbrow <i>q</i> er.cc	ordbrow <i>s</i> er.cc	old + browser

E. Transformer Performance under Severe Corruption

In this section, we demonstrate the Transformer’s success under severe corruption. First, we present examples in which the Transformer successfully decodes heavily corrupted inputs. Recall that in this study, flip errors and burst errors were artificially simulated. Figures 10 and 11 shows cases for both flip and burst errors where the Transformer succeeds in decoding while the pyzbar fails to read the code. In each case, the corruption causes a significant deviation from the original QR code, preventing pyzbar from decoding it. However, the Transformer successfully decodes these corrupted codes. The readers are recommended to try a QR code scan using their smartphone.



Figure 10. Examples where the Transformer successfully decodes QR codes after 20 flip errors. “original” shows the QR code before corruption, and “corrupted” shows the version after applying the errors. While pyzbar fails to decode the corrupted QR code, the Transformer can decode it.



Figure 11. Examples where the Transformer successfully decodes QR codes after 10 burst errors. “original” shows the QR code before corruption, and “corrupted” shows the version after applying the errors. While pyzbar fails to decode the corrupted QR code, the Transformer can decode it.

Next, we examine the similarity between the plain text and the strings generated by the Transformer under severe corruption.

Similarity was computed as the Levenshtein distance (Levenshtein, 1966) normalized by the maximum string length, as follows:

$$\text{Similarity}(a, b) = 1 - \frac{\text{Levenshtein}(a, b)}{\max(|a|, |b|)}. \quad (5)$$

Here, a and b denote the two strings being compared. Figure 12 illustrates the distributions of similarity scores for the generated strings compared to the plain text when subjected to 20 flip errors and 10 burst errors, respectively. In both (a) and (b), the Transformer produces a large proportion of strings with over 80 % similarity to the plain text, even under severe corruption. Moreover, applying data augmentation substantially increases these similarity scores. Therefore, the Transformer can generate highly similar outputs despite heavy damage, and its performance can be further enhanced through data augmentation. For reference, Table 8 presents the correspondence between similarity scores and the actual generated strings.

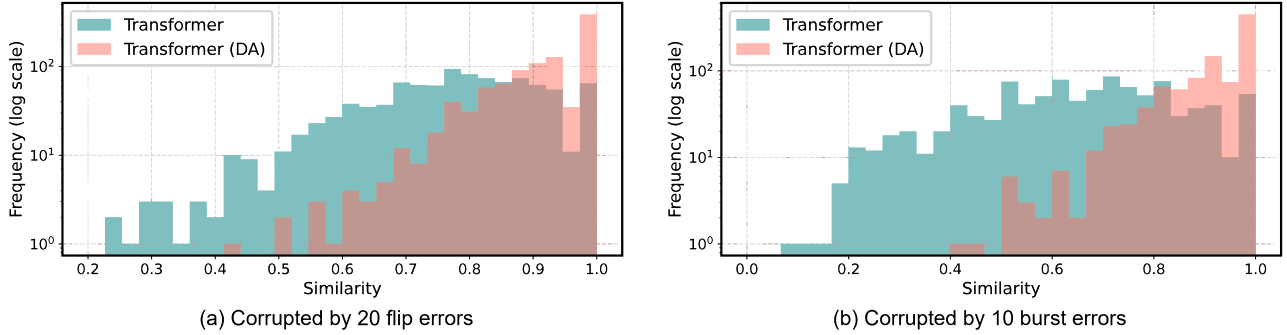


Figure 12. Distributions of similarity scores for the generated strings compared to the plain text when subjected to (a) 20 flip errors and (b) 10 burst errors.

Table 8. Examples of strings and their similarity scores.

Strings	Transformer	Transformey	Transforcel	Transfuimur	Tranyfjjber
Similarity	1.00	0.91	0.82	0.73	0.64
pcavsfzume	buagyfirchr	Trpbgfbildv	xwbiliorblq	bpqzcprwvh	bxxggbggpjx
0.45	0.36	0.27	0.18	0.09	0.00

F. Sensitivity of Decoding QR Code

In this section, we evaluate the sensitivity of QR code decoding. The sensitivity of a function is defined by how much the output changes in response to variations in the input. In this experiment, we examine how many bits in a QR code must be altered to produce a one-character change in the expected output (i.e., the plain text). We prepare a set of domain names and generate 1,000 samples by randomly modifying a single character in each (e.g., changing “example.com” to “exomple.com”). We then convert both the original and modified strings into QR codes. Finally, we count the number of differing bits between the original and modified QR codes and compute the average.

Table 9 shows the average number of bit changes resulting from a one-character change in the plain text for QR code Version 1 through Version 3. For Version 1 and Version 2, the number of bit changes required to alter the output increases with the error correction level. This suggests that the decoding of QR codes in Version 1 and Version 2 becomes less sensitive as the error correction level increases. On the other hand, Version 3 does not follow this trend. This is attributed to differences in the encoding schemes used in each version. In Version 3 and later versions, the data codewords are divided into two or more blocks, and each block is independently encoded with Reed–Solomon codes to generate corresponding error correction codewords. Because Version 3 and later versions perform encoding on split data blocks, the function sensitivity does not necessarily decrease with higher error correction levels. Given these factors, Version 1 and Version 2 are more suitable for examining how function sensitivity impacts the behavior of Transformers.

Table 9. Average number of bit changes caused by a single-character change in the plain text.

Version	L	M	Q	H
1	30.3	42.0	54.1	69.2
2	43.4	66.7	90.2	115.4
3	62.1	106.5	74.3	92.1

G. Sensitivity Analysis Across QR Code’s Four Error-Correction Levels

In this section, we investigate the Transformer’s behavior by training it to decode QR codes at four different error-correction levels, each corresponding to a distinct sensitivity.

Figure 13(a) shows the success rate of the Transformers under flip errors. The results showed that there were almost no differences in success rate between error correction levels. For comparison, Figure 13(b) shows the success rate of Transformer pyzbar. It is clear that robustness against corruption varies significantly with the error-correction level. The performance gap between the Transformer and pyzbar likely arises because the Transformer focuses exclusively on the data codewords that carry the essential plain text information and thus ignores the error correction codewords.

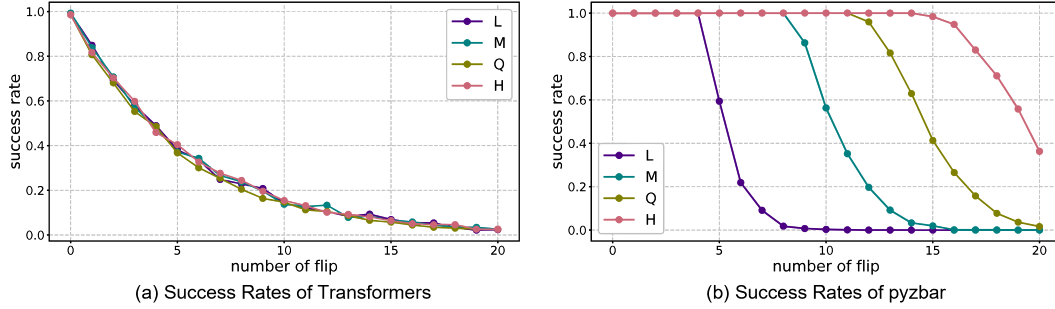


Figure 13. Success rates under flip errors for Transformer (a) and pyzbar (b), evaluated on v2-QR codes with all four error correction levels. The Transformer shows almost no difference in success rates across the error correction levels, whereas the pyzbar exhibits substantial variation depending on the level.

H. Success Rate of Error Correction in the Encoding Region

Here, we derive the success rate of error correction in QR codes with n -bit errors. We assume that a bit flip occurs uniformly at random in the encoding region. The encoding region consists of the data and error-correction codewords, the format information, and the remainder bits. Note that the remainder bits are not used in the decoding; thus, any errors in those bits do not affect the result.

Let N denote the total number of bits in the encoding region, and let N_d , N_f and N_r be the numbers of bits in the data and error-correction codewords and the format information, remainder bits, respectively. Accordingly, the numbers of erroneous bits are denoted by p , q , and $n - p - q$, respectively. Then, the successful rate of error correction is given by the following.

Theorem 4.1 (Success Rate of Error Correction in the Encoding Region). *Let n be the total number of bit errors in the encoding region. Then, the success rate of error correction—denoted by $P_{\text{success}}(n)$ —is given by:*

$$P_{\text{success}}(n) = \frac{1}{\binom{N}{n}} \sum_{p=0}^n \sum_{q=0}^{n-p} W(p, q) P_d(p) P_f(q), \quad (1)$$

where P_d and P_f respectively denote the success rate of error correction in the data and error-correction codewords with p -bit errors and that in format information with q -bit errors, and

$$W(p, q) = \binom{N_d}{p} \binom{N_f}{q} \binom{N_r}{n-p-q} \quad (2)$$

is the probability that $(n - p - q)$ -bit errors fall into the remainder bits.

In the following, we provide an overview of the derivation of $P_d(p)$ and $P_f(q)$.

$P_d(p)$ - Success Rate in Data and Error-Correction Codewords. For the data and error-correction codewords, let M denote the total number of codewords, and let M_{ecc} represent the number of error-correction codewords, which depends on the error correction level. Here, $M = N_d/8$. According to the properties of Reed–Solomon codes, the maximum number of correctable codewords, denoted by t , is determined by $t = \lfloor M_{\text{ecc}}/2 \rfloor$. Under these conditions, the probability of successful error correction for the data and error-correction codewords can be determined as follows:

Theorem 4.2 (Success Rate of Error Correction in Data and Error-Correction Codewords). *Let p be the number of erroneous bits within the data and error-correction codewords. Then, the success rate of error correction—denoted by $P_d(p)$ —is given by*

$$P_d(p) = \frac{1}{\binom{N_d}{p}} \sum_{k=\lceil \frac{p}{8} \rceil}^t \binom{M}{k} \sum_{j=0}^k (-1)^j \binom{k}{j} \binom{8(k-j)}{p}. \quad (3)$$

Proof. Let K denote the number of codewords that contain at least 1-bit error. Since the Reed–Solomon decoder can correct up to t codewords with errors, the Success Rate is given by:

$$\begin{aligned} P_d(p) &= P(K \leq t) \\ &= \sum_{k=\lceil \frac{p}{8} \rceil}^t P(K = k) \\ &= \frac{1}{\binom{N_d}{n}} \sum_{k=\lceil \frac{p}{8} \rceil}^t |\{K = k\}|. \end{aligned} \quad (6)$$

Let $S_k(p)$ denote the number of ways in which exactly k distinct codewords each contain at least 1-bit error, given that there are p bit errors in total. Then, under the condition that all k selected codewords include at least one erroneous bit, the value of $S_k(p)$ is given by

$$|\{K = k\}| = \binom{M}{k} S_k(p). \quad (7)$$

Let U be the set of all ways to choose n error-bit positions out of the $8k$ bits of the k selected codewords (so $|U| = \binom{8k}{p}$). Define $A_i \subseteq U$ as the subset of error assignments in which the i -th codeword contains no bit errors. Then, the number of assignments in which every codeword contains at least 1-bit error is given by:

$$\begin{aligned} S_k(p) &= \left| \bigcap_{i=1}^k A_i^c \right| \\ &= |U| - \left| \bigcup_{i=1}^k A_i \right|. \end{aligned} \quad (8)$$

To evaluate the number of assignments in which every codeword contains at least one erroneous bit, we apply the principle of inclusion-exclusion, as formalized in Lemma I.1. Specifically, applying it to the sets $A_1, A_2, \dots, A_k \subseteq U$, we obtain:

$$\left| \bigcup_{i=1}^k A_i \right| = \sum_{j=1}^k (-1)^{j-1} \sum_{L \subseteq [k], |L|=j} \left| \bigcap_{l \in L} A_l \right| \quad (9)$$

The term $\left| \bigcap_{l \in L} A_l \right|$ denotes the number of cases in which j codewords, selected from the k total codewords, are all error-free; thus,

$$\sum_{L \subseteq [k], |L|=j} \left| \bigcap_{l \in L} A_l \right| = \binom{k}{j} \binom{8(k-j)}{p}. \quad (10)$$

From Equations (8) to (10), it follows that

$$\begin{aligned}
 S_k(p) &= |U| - \left| \bigcup_{i=1}^k A_i \right| \\
 &= \binom{8k}{p} - \sum_{j=1}^k (-1)^{j-1} \binom{k}{j} \binom{8(k-j)}{p} \\
 &= \binom{8k}{p} + \sum_{j=1}^k (-1)^j \binom{k}{j} \binom{8(k-j)}{p} \\
 &= \sum_{j=0}^k (-1)^j \binom{k}{j} \binom{8(k-j)}{p}.
 \end{aligned} \tag{11}$$

Therefore, from Equations (6), (7) and (11), we obtain

$$\begin{aligned}
 P_d(p) &= \frac{1}{\binom{N_d}{p}} \sum_{k=\lceil \frac{p}{8} \rceil}^t |K = k| \\
 &= \frac{1}{\binom{N_d}{p}} \sum_{k=\lceil \frac{p}{8} \rceil}^t \binom{M}{k} \sum_{j=0}^k (-1)^j \binom{k}{j} \binom{8(k-j)}{p}.
 \end{aligned} \tag{12}$$

□

$P_f(q)$ - Success Rate in Format Information. The format information is encoded into 15 bits using a (15, 5) BCH code and placed in two separate locations in the QR code. If error correction works in either one of the two, the format information can be correctly read. The (15, 5) BCH code can correct up to 3-bit errors. Under these conditions, the probability that the format information is correctly recovered is given below:

Theorem 4.3 (Success Rate of Error Correction in Format Information). *Let q denote the number of erroneous bits affecting the format information, and let i and j represent the number of bit errors in each of the two respective instances. Then, the success rate of error correction—denoted by P_f —is given by*

$$P_f(q) = \frac{|\{(i, j) \in \mathbb{N}_0^2 \mid i + j = q, \min(i, j) \leq 3\}|}{q + 1}. \tag{4}$$

Proof. When exactly q bit-errors occur in the format information, there are $q + 1$ ways to distribute those errors across the two (15, 5) BCH codeword blocks. Decoding of the format information succeeds if at least one of the two blocks contains at most 3-bit errors. Labeling by i and j the numbers of errors in the first and second blocks respectively, the number of allocations that lead to successful decoding can be written as $|\{(i, j) \in \mathbb{N}_0^2 \mid i + j = q, \min(i, j) \leq 3\}|$. Hence,

$$P_f(q) = \frac{|\{(i, j) \in \mathbb{N}_0^2 \mid i + j = q, \min(i, j) \leq 3\}|}{q + 1}. \tag{13}$$

□

I. Principle of Inclusion–Exclusion

Lemma I.1 (Principle of Inclusion–Exclusion). *Let A_1, A_2, \dots, A_k be subsets of a finite set U . Then, the cardinality of their union is given by:*

$$\left| \bigcup_{i=1}^k A_i \right| = \sum_{j=1}^k (-1)^{j-1} \sum_{L \subseteq [k], |L|=j} \left| \bigcap_{l \in L} A_l \right|, \tag{14}$$

where $[k] = \{1, 2, \dots, k\}$.

Proof. We prove the identity by counting how many times each element $x \in U$ is counted on the right-hand side. Fix an element $x \in U$, and suppose that x belongs to exactly m of the sets A_1, A_2, \dots, A_k . Without loss of generality, assume $x \in A_1, A_2, \dots, A_m$, and $x \notin A_{m+1}, \dots, A_k$. On the right-hand side, x contributes to each term of the form $|\bigcap_{l \in L} A_l|$, where $L \subseteq [k]$ and $x \in A_l$ for all $l \in L$. Since x belongs to exactly m of the A_i , the total number of such contributing terms is:

$$(-1)^0 \binom{m}{1} + (-1)^1 \binom{m}{2} + \dots + (-1)^{m-1} \binom{m}{m} = \sum_{j=1}^m (-1)^{j-1} \binom{m}{j}. \quad (15)$$

This sum can be evaluated using the binomial theorem:

$$\begin{aligned} 0 &= (1 - 1)^m \\ &= \sum_{j=0}^m (-1)^j \binom{m}{j} \\ &= 1 - \sum_{j=1}^m (-1)^{j-1} \binom{m}{j}. \end{aligned} \quad (16)$$

From Equation (16), it follows that

$$\sum_{j=1}^m (-1)^{j-1} \binom{m}{j} = 1. \quad (17)$$

Therefore, every $x \in \bigcup_{i=1}^k A_i$ contributes exactly once to the right-hand side. Elements not in the union contribute zero. Hence, the right-hand side counts exactly the number of elements in the union. \square