



End-to-end semantic preservation in text-aware image compression systems^{☆,☆☆}

Stefano Della Fiore^{*}, Alessandro Gnutti, Marco Dalai, Pierangelo Migliorati, Riccardo Leonardi

Department of Information Engineering, University of Brescia, Italy

ARTICLE INFO

Keywords:

End-to-end image compression
Coding for machines

ABSTRACT

Traditional image compression methods aim to reconstruct images for human perception, prioritizing visual fidelity over task relevance. In contrast, Coding for Machines focuses on preserving information essential for automated understanding. Building on this principle, we present an end-to-end compression framework that retains text-specific features for Optical Character Recognition (OCR). The encoder operates at roughly half the computational cost of the OCR module, making it suitable for resource-limited devices. When on-device OCR is infeasible, images can be efficiently compressed and later decoded to recover textual content. Experiments show significant improvements in text extraction accuracy at low bitrates, even outperforming OCR on uncompressed images. Building on these insights, we further explore general-purpose encoders under extreme compression, investigating whether compact, visually degraded representations can still retain recoverable semantic information. Results demonstrate that semantic content can persist despite severe compression, bridging task-oriented text compression and broader machine-centered image coding.

1. Introduction

For decades, image compression has been a crucial research area, driven by the increasing demand for efficient data storage and transmission across diverse digital applications. Alternative representations based on graph signal processing and graph Fourier transforms have also been explored for image and residual coding [2–6]. The advent of deep learning has transformed this field, driving significant interest in end-to-end learned compression frameworks, with variational autoencoder (VAE)-based methods playing a key role [7,8]. Unlike traditional techniques, these approaches jointly optimize the encoding and decoding stages within a unified framework. Recent works [9–14] have demonstrated that learned compression techniques can outperform VVC [15], the state-of-the-art standard for image and video coding, in terms of key quality metrics such as Peak Signal-to-Noise Ratio (PSNR) and Multi-Scale Structural Similarity (MS-SSIM). These findings underscore the transformative potential of learned approaches in shaping the future of image coding [16].

Although most learned image compression methods are tailored for human perception, the growing prominence of machine vision

applications has driven a shift toward task-specific optimization. Recently, image coding for machine perception has gained significant attention, fueled by the increasing need to transmit visual data efficiently for high-level recognition tasks across devices. Referred to as Image Coding for Machines (ICM), differently from general image enhancement methods [17,18], this strategy aims to refine the compression process to generate images better suited for specific downstream tasks. Some examples of tasks in this domain include denoising [19], deepfake detection [20], computer vision tasks [21], face detection in the compressed domain [22,23], and, more recently, Multimodal Large Language Models (MLLMs) [24].

Among these recognition applications, Optical Character Recognition (OCR) is a technology that enables the automatic extraction of text from images, scanned documents, and other visual data sources. OCR systems can convert printed or handwritten text into machine-readable formats. This capability is widely used across various applications, including document digitization, automated data entry, and assistive technologies for the visually impaired. Modern OCR systems often employ deep neural networks to enhance accuracy, enabling them to handle diverse fonts, handwriting styles, and challenging conditions such as poor lighting or distortions.

[☆] This article is part of a Special issue entitled: 'IMAGE_Visual Data Coding' published in Signal Processing: Image Communication.

^{☆☆} This work was supported in part by the European Union under the Italian National Recovery and Resilience Plan (NRRP) of NextGenerationEU, Partnership on "Telecommunications of the Future", Program "RESTART" under Grant PE00000001, "Netwin" Project (CUP E83C22004640001) and "FRAME" Project (CUP C89J24000250004). An earlier version of this paper was presented in part in [1] at EUSIPCO 2025.

^{*} Corresponding author.

E-mail address: stefano.dellafiore@unibs.it (S. Della Fiore).




Fig. 1. Comparison of high-level frameworks: no compression, conventional compression, and our proposed TFIC.

However, a major challenge arises when applying image compression to OCR-relevant images. Compression techniques, especially at low bitrates, can introduce significant distortions, such as blurring, blocking artifacts, and loss of fine details, which can severely impact text readability. Since OCR performance heavily depends on the clarity and integrity of textual structures, excessive compression may lead to character misinterpretation or complete recognition failure.

In this paper, we first introduce an image compression system designed specifically for OCR-based vision tasks. Our learnt Text-Focused Image Compression (TFIC) model (see Fig. 1) is built on a standard architecture commonly used for neural image compression. However, we incorporate an OCR model at the system's output and optimize the compression process by minimizing the loss specific to OCR performance. This approach allows us to achieve a balance between efficient compression and high-quality text recognition, even at extremely lower compression rates where traditional methods strongly degrade the text to the point of being unreadable.

Then, we further extend this study to general-purpose encoders, exploring their capacity to preserve hidden semantics under extreme compression. Instead of optimizing for visual fidelity, we examine whether compact, visually degraded representations can retain recoverable meaning through learned enhancement and recognition modules. Indeed, while traditional compression systems, such as JPEG, are designed to minimize perceptual distortion at typical operating points, they are not tailored to preserve semantic content at extremely low bitrates. In scenarios such as large-scale data transmission, where bandwidth is critically limited, operating at these extreme compression levels may be unavoidable. The main challenge is that such aggressive quantization typically removes most of the low-level features necessary for human or machine interpretation. As a result, the reconstructed images appear as noisy, blurred, or completely corrupted patterns.

In the second part of our work, we thus present a post-decoder enhancement system that leverages prior knowledge through deep learning models to recover semantic information, even from visually degraded signals. In particular, we focus on text, showing that semantic content can be implicitly preserved and successfully extracted despite severe compression.

Thus, while some recent works aim at improving semantic representation through advanced feature modeling strategies, such as graph-based or frequency-aware interactions [25,26], our work focuses on task-aware image compression and, to the best of our knowledge, is among the first to explicitly optimize the compression process for OCR-driven semantic preservation under strong bitrate constraints.

The remainder of the paper is organized as follows. Section 2 provides an overview of neural image compression and OCR systems. Then, the first part of the paper (Section 3) presents our TFIC model for OCR tasks, describing its architecture, end-to-end training, and experimental results with comparisons to baseline methods. Finally, in the second part (Section 4), we extend our analysis to the use of general-purpose image encoders, investigating their behavior outside conventional operating conditions.

2. Background

2.1. Neural image compression

An end-to-end learned image compression system generally consists of two key components: the main autoencoder and the hyperprior autoencoder. The main autoencoder is composed of an analysis transform g_a and a synthesis transform g_s . The analysis transform g_a encodes an RGB image $x \in \mathbb{R}^{H \times W \times 3}$, with height H and width W , into a latent representation $y \in \mathbb{R}^{\frac{H}{16} \times \frac{W}{16} \times C_y}$ using an encoding distribution $q_{g_a}(y|x)$. This latent y is then uniformly quantized as $\hat{y}$ and entropy encoded into a bitstream via a learned prior distribution $p(\hat{y})$. On the decoder side, $\hat{y}$ is entropy decoded and reconstructed as $\hat{x} \in \mathbb{R}^{H \times W \times 3}$ using a decoding distribution $q_{g_s}(\hat{x}|\hat{y})$, implemented by the synthesis transform g_s . Throughout this process, the prior distribution $p(\hat{y})$ plays a critical role in determining the number of bits needed to represent the quantized latent $\hat{y}$. To mitigate this, a hyperprior autoencoder [8] is introduced to model the prior distribution in a content-adaptive manner. The hyperprior autoencoder consists of a hyperprior analysis transform h_a and a hyperprior synthesis transform h_s . The transform h_a converts the image latent y into side information $z \in \mathbb{R}^{\frac{H}{64} \times \frac{W}{64} \times C_z}$, which typically represents a small portion of the compressed bitstream. This quantized version of z is then decoded from the bitstream through h_s , resulting in the learned prior distribution $p(\hat{y})$.

2.2. OCR systems

An OCR system is designed to extract text from images captured in unconstrained environments. The system generally consists of four integrated modules. First, a *detection* module identifies and localizes regions containing text within an image, often using object detection frameworks to output bounding boxes around potential text areas [27–29]. Next, the *transformation* stage normalizes these detected regions to correct for distortions such as skew, rotation, or perspective changes. For instance, a Spatial Transformer Network (STN) with a Thin-Plate Spline (TPS) [30] transformation may be employed to obtain a rectified image, denoted as $\tilde{x}$.

Subsequently, a *feature extraction* component, typically implemented with a convolutional neural network (CNN) like ResNet [31], converts the (transformed) image x or $\tilde{x}$ into a rich feature map $V = \{v_i\}_{i=1}^I$, where each v_i represents features corresponding to a specific region of the image and I is the dimension of the feature map. Given the sequential nature of text, these features are then reorganized into a sequence and processed by a *sequence modeling* module, commonly a bidirectional LSTM (BiLSTM) (see [32–34]) or Transformer, that captures contextual dependencies between characters, yielding a contextualized representation. Finally, the *prediction* stage decodes this representation into the final text output using methods such as Connectionist Temporal




Fig. 2. High-level architectural framework of TFIC. The main image codec is composed by an encoder g_a that compresses the input image into a latent representation, which is quantized and transmitted. The decoder g_s reconstructs the image that is subsequently processed by the OCR module for text extraction.

Classification (CTC) [35] or attention-based decoding [33,34]. The entire pipeline is trained end-to-end by minimizing a loss function that measures the discrepancy between the predicted sequence and the ground truth, ensuring robust performance in complex and real-world scenarios.

3. Proposed text-focused image compression

3.1. Architecture overview

Fig. 2 illustrates our proposed framework. The image codec architecture follows the key modules outlined in 2.1. As a reference, we adopt the image compression model from [36], which utilizes Transformer-based backbones to implement both the main and hyper-prior autoencoders.

An OCR system with frozen parameters is integrated downstream of the decoder alongside the autoencoder. Its architecture follows the key modules outlined in 2.2, with the Scene Text Recognition model from [37] serving as a reference. During training, the OCR extracts text from $\hat{x}$, denoted as $T(\hat{x})$, which is compared to the ground truth text $G(x)$. The resulting loss signal is then backpropagated through the entire network. This process helps guide both the encoder and decoder to preserve text-relevant information.

3.2. Loss function and end-to-end training

The overall training objective $\mathcal{L}_{\text{total}}$ of our method is a combination of three loss components:

$$\mathcal{L}_{\text{total}} = \lambda \cdot \mathcal{L}_{\text{dist}}(x, \hat{x}) + \mathcal{L}_{\text{rate}}(\hat{y}, \hat{z}) + \gamma \cdot \mathcal{L}_{\text{OCR}}(G(x), T(\hat{x})), \quad (1)$$

where:

- $\mathcal{L}_{\text{dist}}(x, \hat{x})$ is the distortion loss, measured by the mean squared error (MSE) between the original image x and its reconstruction $\hat{x}$.
- $\mathcal{L}_{\text{rate}}(\hat{y}, \hat{z}) = -\log p(\hat{z}) - \log p(\hat{y}|\hat{z})$ represents the rate loss, which estimates the bitstream length needed to encode the quantized latent representation $\hat{y}$ and side information $\hat{z}$.
- $\mathcal{L}_{\text{OCR}}(G(x), T(\hat{x}))$ is the OCR loss that quantifies the difference between the ground truth text $G(x)$ and the text recognized from the reconstructed image $T(\hat{x})$. This is measured using the cross-entropy loss between the text-indexes of the ground truth text and the reconstructed one.

The parameters λ and γ control the trade-off between pixel-level fidelity and text recognition accuracy. In particular, increasing λ in Eq. (1) favors visually faithful reconstructions, while increasing γ promotes the preservation of features that are most relevant for text recognition. In this work, we are primarily interested in the semantic operating regime, where the objective is to maximize OCR performance rather than visual fidelity. This corresponds to prioritizing the OCR loss

(i.e., $\lambda = 0$). However, directly optimizing this objective from scratch leads to unstable training and poor convergence. For this reason, we adopt a two-stage training strategy. In the first stage, we set $\gamma = 0$, allowing the model to be pre-trained using only the distortion and rate terms, with the target rate controlled by λ , following the standard practice in learned image compression to obtain models operating at different rate-distortion trade-offs. This stage enables the model to learn a stable and meaningful image representation. In the second stage, we set $\lambda = 0$ and optimize only the OCR loss, with γ fixed to 0.1 for all rate points. This shifts the optimization toward semantic preservation, encouraging the model to retain only the information that is relevant for text recognition.

Training is conducted end-to-end, enabling the encoder, decoder to co-adapt such that the final compressed representation retains essential textual information while achieving high compression efficiency. We recall that only the image codec is re-trained, while the OCR system remains unchanged.

3.3. Dataset and experimental setup

We use a synthetic data generator for text recognition [38] to create training and test sets of approximately 20k and 600 images, respectively. Each image is resized to a fixed resolution of 256×512 pixels. The dataset covers a diverse range of fonts, layouts, and background complexities, ensuring a comprehensive evaluation of text extraction under various conditions.

For evaluation, we compare the proposed method with the base codec optimized exclusively for MSE on the same training set, thus without including the OCR loss component. Both models are evaluated based on bitrate (bits-per-pixel, bpp) and OCR performance, which is measured by the Levenshtein edit-distance $\text{lev}(G(x), T(\hat{x}))$ (see [39]) that represents the minimum number of single-character edits (insertions, deletions or substitutions) required to change the reconstructed text $T(\hat{x})$ into the ground truth one $G(x)$. The accuracy is then defined as

$$1 - \frac{\text{lev}(G(x), T(\hat{x}))}{\max\{|G(x)|, |T(\hat{x})|\}}, \quad (2)$$

where $|G(x)|$ and $|T(\hat{x})|$ are the number of alphanumeric characters contained in the ground truth text and in the reconstructed one, respectively.

3.4. Results

Fig. 3(a) shows the rate vs. OCR performance curves for both the base codec and the proposed TFIC model. The base codec exhibits a significant drop in OCR accuracy defined in Eq. (2), with lower accuracy values at comparable bitrates. In contrast, our method consistently achieves higher accuracy, effectively preserving text information even at lower bitrates. Notably, TFIC even surpasses the OCR performance on uncompressed images, suggesting that it acts as a beneficial pre-processing step. This behavior can be attributed to the fact that the




Fig. 3. Comparison of OCR and PSNR performance across different methods.




Fig. 4. Visual comparison of reconstructed images. Although the base codec preserves more global details, our method retains critical textual regions that yield superior OCR performance.

proposed OCR-guided compression does not aim to preserve the full visual content of the image, but rather to retain and enhance features that are most relevant for text recognition. As a result, background details, textures, and other non-textual elements that may act as noise for the OCR model are suppressed. This leads to simplified representations in which text regions are emphasized while irrelevant content is discarded, effectively providing the OCR module with cleaner and more discriminative inputs. In fact, while the base codec achieves higher overall PSNR (see Fig. 3(b)), text regions often appear blurred or distorted, resulting in poor OCR accuracy. In contrast, our method preserves sharp and legible text, even if non-essential areas of the image undergo destructive compression.

A comparison of the results shown in the second, third and fourth columns of Fig. 4, shows that our method can also be useful when used *in addition* to a standard compression scheme, whenever transmission of a low rate visual content of the background is also required. Indeed, the total rate for sending *both* images in the third and fourth columns (as two separate layers) equals the rate for the image in the second column.

The computation time would in this case be doubled, but would still remain lower than the time needed for the OCR on the original image alone, while also doing compression.

3.5. Failure cases and limitations

Although the proposed TFIC model achieves strong average OCR performance at very low bitrates, its behavior is not uniformly robust across all text-image configurations. Fig. 5 reports three representative failure cases observed in our experiments.

The first two examples highlight a challenging scenario in which thin characters are overlaid on highly textured urban backgrounds. In these cases, background edges and text strokes exhibit similar spatial frequencies and comparable local contrast. Under severe bitrate constraints, the encoder cannot always separate the most informative textual structures from the surrounding clutter, and the reconstructed image tends to retain only fragmented or distorted character traces. As




Fig. 5. Representative failure cases for the proposed TFIC model. In challenging scenarios involving thin characters, low-contrast text, or highly structured backgrounds, the reconstructed image may preserve only partial stroke information, leading to merged characters or severe loss of legibility.

a result, OCR may fail because character boundaries become ambiguous or partially missing.

The third example shows a different failure mode. Here, the word is long, weakly contrasted, and placed over a repetitive brick pattern. The periodic texture of the background competes with the text structure, while the low contrast reduces the saliency of the characters. In this situation, the semantic representation learned by TFIC may allocate too few bits to the textual component, producing a reconstruction in which the word is almost completely suppressed.

These examples point to some limitations of the proposed method. First, TFIC is more effective when text regions remain sufficiently distinguishable from the background, whereas performance degrades for thin fonts, reduced inter-character spacing, low-contrast text, or strong structured interference. Second, because the model is optimized mainly through the OCR objective, it may preserve only the cues that are statistically useful for recognition on average, without guaranteeing a faithful reconstruction of every character in particularly difficult scenes. Finally, the synthetic training setup, although diverse, may not fully cover challenging cases such as transparent text overlays, strongly cluttered natural backgrounds, or very long words with weak stroke continuity.

Future improvements could therefore include harder training samples with more complex backgrounds, explicit text localization or attention mechanisms, and auxiliary losses promoting stroke continuity and character separation in addition to OCR accuracy.

3.6. Runtime analysis

One of the motivations behind our method is the difference in computational cost between the encoding and OCR stages. Empirical measurements show that the encoding stage requires only about half

Table 1

Average and standard deviation of the encoding and OCR processing times per image, together with the number of model parameters and the computational cost in FLOPs per pixel.

	Encoding	OCR module
Computational time [ms]	12.9 ± 1.8	24.1 ± 3.3
Parameters	$7.1 \cdot 10^6$	$47.9 \cdot 10^6$
FLOPs/pixel	$1.04 \cdot 10^3$	$3.29 \cdot 10^6$

of the time needed to perform OCR on the same image. This is particularly relevant for resource-constrained devices, where fast on-device compression is desirable, while the more demanding OCR stage can be executed later on external hardware.

Table 1 summarizes the computational characteristics of the two modules. In addition to the average processing time per image, reported with standard deviation, the table also includes the number of model parameters and the computational cost expressed in FLOPs per pixel. The results show that the encoding stage is not only faster at runtime, but also substantially lighter than the OCR module in terms of both model size and computational complexity. This further supports the practicality of the proposed framework, where lightweight on-device compression can be performed efficiently, while the more demanding OCR processing can be deferred or executed remotely.

4. Hidden semantic preserved by general-purpose encoders

In the second phase of our study, we expanded our investigation to the use of general-purpose encoders, particularly focusing on how these encoders behave outside their standard operating ranges when the objective is not mere visual fidelity but rather the preservation of




Fig. 6. Comparison of high-level frameworks: no compression, regular compression, and our proposed method.

semantic content. The underlying question is whether it is possible to exploit extremely compressed representations, although corrupted and unintelligible in a conventional sense, recovering meaningful semantic information through learned enhancement and recognition modules.

4.1. Case study: Textual images

As a case use, we focused on textual images subjected to extremely low-rate JPEG compression. The dataset consisted of low-resolution digital text images and these images were compressed using JPEG with extreme quantization levels, leading to a representation that to the human eye resembles random noise or blotches of black and white pixels. Unsurprisingly, applying a conventional OCR system to these compressed images yields recognition rates close to zero, confirming that no usable information can be extracted through standard pipelines.

Despite this, we hypothesized that the statistical correlations present in such degraded images could still retain implicit traces of the underlying semantic content. To exploit this, we trained convolutional neural networks specifically designed to map compressed JPEG inputs back to a legible textual domain.

In Fig. 6, a high-level diagram of our method is presented. The block JPEG LR-LQ represents a module that performs JPEG compression on low-resolution images using extremely high quantization, leading to significant compression and loss of detail. The block DEQUANT denotes a convolutional module designed to perform dequantization, effectively learning to restore image details that have been degraded by the high quantization in the JPEG compression stage.

The first step involved designing a convolutional neural network to post-process the decoded JPEG images. Even with a simple mean squared error (MMSE) objective function, the network was able to restore structural features of the characters, producing outputs sufficiently clear for OCR modules to operate. This result highlights that even when low-level image fidelity is lost, mid-level structures associated with characters remain recoverable.

The architecture adopted in our preliminary experiments consisted of a lightweight convolutional neural network with residual connections, allowing the model to focus on refining corrupted strokes and suppressing JPEG artifacts. Interestingly, the restored images did not always visually resemble the originals in terms of fine details; however, they were interpretable at a semantic level, sufficient for both OCR recognition and human readability.

4.2. Architecture overview

The proposed network is based on a modified DnCNN architecture designed to perform image restoration and dequantization on JPEG-compressed low-resolution patches. As illustrated in Fig. 7, the model follows a fully convolutional design composed of a series of feature extraction, normalization, and residual learning blocks.




Fig. 7. High-level schematic representation of the DEQUANT block that is a DnCNN-based architecture used for image dequantization. The model predicts the residual component, which is subtracted from the input to obtain the restored image.

The network receives as input an image patch of size 64×256 and processes it through an initial convolutional layer with 64 filters of size 3×3 , followed by a ReLU activation. Subsequently, a sequence of $(D - 2)$ convolutional blocks, each composed of a 3×3 convolution, batch normalization (BN), and ReLU activation, progressively refines the feature representations. In our implementation, we set the network depth to $D = 17$, resulting in 15 intermediate convolutional layers with batch normalization. The final convolutional layer produces a residual output $\hat{r}$ of the same spatial dimensions as the input.

The network is trained to estimate the residual component (i.e., the quantization noise) rather than the clean image directly. The restored image $\hat{x}$ is obtained as

$$\hat{x} = x - \hat{r}, \quad (3)$$

where x denotes the input patch and $\hat{r}$ is the predicted noise component.

4.3. Loss function

The training objective is defined as a combination of a pixel-domain reconstruction loss and a frequency-domain consistency loss. The overall loss function is expressed as

$$\mathcal{L}_{\text{total}} = \lambda \cdot \mathcal{L}_{\text{dist}}(x, \hat{x}) + (1 - \lambda) \cdot \mathcal{L}_{\text{DCT}}(x, \hat{x}), \quad (4)$$

where $\lambda \in [0, 1]$ is a weighting coefficient that balance the contribution of each term.

The first term in Eq. (4), $\mathcal{L}_{\text{dist}}(x, \hat{x})$, corresponds to the classical mean squared error loss between the reconstructed image $\hat{x}$ and the original image x which enforces pixel-wise fidelity in the spatial domain.




Fig. 8. Schematic representation of the proposed OCR-guided dequantization architecture (DEQUANT block). A frozen DnCNN produces an initial reconstruction, which is analyzed by a CRAFT detector to extract word-level patches. Each patch is refined by a trainable DnCNN and evaluated by a fixed OCR module that provides a recognition-based loss for training.

The second term, $\mathcal{L}_{\text{DCT}}(x, \hat{x})$, is a blockwise Discrete Cosine Transform (DCT) loss designed to enforce frequency-domain consistency between the reconstructed image $\hat{x}$ and the input quantized image x . Let B denote the set of non-overlapping 8×8 blocks extracted from the images x and $\hat{x}$. For each block $b \in B$, let $x_b, \hat{x}_b \in \mathbb{R}^{8 \times 8}$ denote the corresponding image patches. Then, the two-dimensional DCT is applied to each block, following the standard JPEG formulation. The resulting DCT coefficients are compared through a thresholded absolute difference:

$$\mathcal{L}_{\text{DCT}}(x, \hat{x}) = \mathbb{E}_{b \sim B} [\max(|\text{DCT}(x_b) - \text{DCT}(\hat{x}_b)| - \tau, 0)], \quad (5)$$

where $\tau = 0.5$ is a fixed margin that suppresses small, perceptually irrelevant differences. This formulation helps the reconstructed image to remain coherent with the quantized input in the frequency domain, ensuring that the reconstructed output does not deviate from the spectral structure imposed by the JPEG compression process.

The combination of $\mathcal{L}_{\text{dist}}$ and $\mathcal{L}_{\text{DCT}}$ enables the network to jointly optimize for spatial accuracy and frequency-domain consistency, leading to reconstructions that are both visually faithful and perceptually coherent.

After a preliminary analysis, we observed that the best reconstruction quality was achieved by setting the weighting coefficient to $\lambda = 0.8$.

4.4. Extended architecture with OCR-guided loss

To further enhance the perceptual and semantic quality of the reconstructed text regions, we extend the baseline DnCNN-based model by introducing an additional *OCR-guided loss* through an integrated text understanding pipeline. The proposed architecture is schematically illustrated in Fig. 8.

Given an input degraded image x (affected by compression or quantization artifacts), the first stage consists of a frozen DnCNN network D_1 , pre-trained according to the method described in Section 4.2. This module produces an initial reconstruction $\tilde{x} = D_1(x)$, which serves as an improved visual prior for the subsequent text analysis.

The reconstructed image $\tilde{x}$ is then processed by a frozen CRAFT (Character Region Awareness for Text Detection) network [27] which identifies and localizes individual word regions. For each detected text region, a corresponding image patch $\tilde{x}_p$ is extracted, forming the set of word-level patches $\mathcal{P}$.

Each patch $\tilde{x}_p$ is subsequently fed into a trainable DnCNN module D_2 , which refines the local reconstruction:

$$\hat{x}_p = D_2(\tilde{x}_p). \quad (6)$$

The refined patch $\hat{x}_p$ is then passed to a fixed OCR recognition network placed at the end of the processing chain. Although the weights of the OCR module remain frozen during training, the backpropagation of the loss gradients goes through the OCR module, enabling D_2 to be optimized according to the semantic feedback produced by the OCR.

An OCR-based loss term $\mathcal{L}_{\text{OCR}}$ is introduced, defined as the discrepancy between the OCR predictions on the reconstructed patches and their corresponding ground-truth text labels $G_p(x)$'s.

$$\mathcal{L}_{\text{OCR}} = \mathbb{E}_{p \sim \mathcal{P}} [\mathcal{L}_{\text{OCR}}(\text{OCR}(D_2(\tilde{x}_p)), G_p(x))], \quad (7)$$

where the OCR-based loss $\mathcal{L}_{\text{OCR}}$ is defined as the cross-entropy between the OCR predictions and the ground-truth textual labels of each detected word.

Finally, all reconstructed patches $\hat{x}_p$ are reassembled to obtain the final restored image $\hat{x}$.

The training phase of the proposed OCR-guided architecture was conducted starting from the pretrained DnCNN model described in Section 4.2. During this stage, all the parameters of the pretrained D_1 , the CRAFT detector, and the OCR recognition module were kept frozen, while only the weights of the trainable D_2 were updated. Fine-tuning was performed exclusively using the OCR-based loss $\mathcal{L}_{\text{OCR}}$. This strategy allows the network to adapt its reconstruction process to maximize the text recognition accuracy, refining the visual details specifically in regions containing textual information.

4.5. Dataset and experimental results

The dataset used for the experiments consists of text images extracted from printed books and documents. Each image contains black text rendered in different fonts and typographic styles on a uniform white background, reproducing realistic scanned document conditions. In total, 2285 full-page images were collected.

To simulate realistic degradation, each clean image was first down-scaled to a lower spatial resolution and then compressed using the JPEG codec with the highest quantization factor, corresponding to extremely aggressive compression. This process generates low-quality, artifact-rich images characterized by severe blocking and quantization effects, which represent the degraded inputs used for network training.

To increase the dataset diversity and sample count, from each original image we randomly extracted 10 non-overlapping patches of fixed size 64×256 pixels. This procedure yields a total of 22850 image patches, each paired with its corresponding quantized version.

The dataset was subsequently divided into two disjoint subsets: 80% of the patches were used for training, while the remaining 20% were reserved for testing and validation.




Fig. 9. Visual results obtained using the baseline DnCNN architecture. The model successfully removes quantization artifacts and reconstructs most visual details, although some text regions remain imperfectly restored. The OCR accuracy achieved by this model was approximately 75%.




Fig. 10. Comparison between the baseline DnCNN and the proposed OCR-guided architecture. The baseline DnCNN achieved an OCR accuracy of approximately 75%, while the proposed OCR-guided model reached about 90%.

In the following, we present the experimental results obtained with the proposed architectures. We first evaluate the baseline model, based on a single DnCNN network trained with the combined MSE and DCT losses described in Section 4.2. Then, we show the improvements achieved by the extended OCR-guided architecture introduced in Section 4.4.

The first set of experiments aims to assess the reconstruction quality achieved by the baseline DnCNN model operating directly on JPEG compressed and quantized images. Fig. 9 shows some qualitative results. Each example shows, from left to right: the degraded input image, the corresponding clean ground-truth, and the restored output produced by the baseline DnCNN. As can be observed, the network effectively removes most of the compression artifacts and restores fine details such as edges and text contours. Although the reconstructions are visually accurate and preserve global image structure, some residual degradation can still be noticed where minor distortions or character ambiguities may remain.

To address the remaining inconsistencies observed in textual areas, we further evaluated the proposed OCR-guided architecture described in Section 4.4. This extended model integrates semantic feedback from a frozen OCR module during training, allowing it to refine the reconstruction.

As shown in Fig. 10, several qualitative differences can be observed between the reconstructions obtained using the baseline DnCNN and those produced by the OCR-guided network. In particular, some words reconstructed by the baseline DnCNN exhibit incorrect or malformed characters, resulting in textual sequences that do not correspond to valid words or that differ semantically from the ground truth. In contrast, the reconstructions obtained with the proposed OCR-guided approach demonstrate an improvement in textual consistency.

In addition to the qualitative analysis, we evaluated OCR accuracy on a test set composed of 4570 patches, each of size 64×256 (Table 2). The OCR accuracy on the original quantized and heavily compressed images was extremely low, remaining below 5%, which confirms the

Table 2OCR accuracy on a test set of 4570 patches of size 64×256 .

Method	OCR accuracy (%)
Original quantized/heavily compressed	< 5
Pureformer [40]	< 10
Baseline DnCNN	75
Proposed OCR-guided architecture	90

severe degradation introduced by the strong quantization. We also considered Pureformer [40], a recent transformer-based image denoising method, as an additional baseline. However, its OCR accuracy remained below 10%, suggesting that a general-purpose denoising approach is not sufficient to recover the fine-grained textual details required for reliable recognition in this setting. When applying the baseline DnCNN model, the accuracy increased substantially, reaching approximately 75%, indicating that the network was able to restore most of the visual information necessary for OCR. Finally, the proposed OCR-guided architecture achieved approximately 90% accuracy, demonstrating that the inclusion of the OCR-based loss helps the model refine text details in a semantically consistent way, leading to significant improvements in both readability and recognition reliability.

5. Conclusions

In this paper, we presented a study on machine-centered image compression with a focus on text and semantic preservation. The work is structured in two complementary parts. In the first part, we introduced a Text-Focused Image Compression (TFIC) system tailored for OCR tasks, showing that integrating task-specific knowledge directly into the compression process enables efficient storage and high text recognition accuracy, even at very low bitrates. In the second part, we explored general-purpose image encoders under extreme compression and proposed post-decoder enhancement systems that leverage deep learning to recover semantic information. Our results demonstrate that meaningful content, particularly textual semantics, can persist despite severe visual degradation, bridging task-oriented compression and more general machine-centered representations.

CRedit authorship contribution statement

Stefano Della Fiore: Writing – review & editing, Writing – original draft, Validation, Software, Conceptualization. **Alessandro Gnutti:** Writing – review & editing, Writing – original draft, Methodology, Investigation, Formal analysis, Conceptualization. **Marco Dalai:** Writing – review & editing, Writing – original draft, Supervision, Methodology, Investigation, Formal analysis, Conceptualization. **Pierangelo Migliorati:** Supervision, Project administration, Funding acquisition. **Riccardo Leonardi:** Supervision, Project administration, Investigation, Funding acquisition.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Data availability

Data will be made available on request.

References

- [1] S. Della Fiore, A. Gnutti, M. Dalai, P. Migliorati, R. Leonardi, TFIC: End-to-end text-focused image compression for coding for machines, in: 2025 33rd European Signal Processing Conference, EUSIPCO, 2025, pp. 1382–1386.
- [2] A. Gnutti, F. Guerrini, R. Leonardi, A. Ortega, Symmetry-based graph Fourier transforms for image representation, in: Proc. 25th IEEE Int. Conf. on Image Proc. (ICIP), 2018, pp. 2575–2579, <http://dx.doi.org/10.1109/ICIP.2018.8451434>.
- [3] A. Gnutti, F. Guerrini, R. Leonardi, A. Ortega, Coding of image intra prediction residuals using symmetric graphs, in: Proc. 26th IEEE Int. Conf. on Image Proc. (ICIP), 2019, pp. 131–135, <http://dx.doi.org/10.1109/ICIP.2019.8803826>.
- [4] A. Gnutti, F. Guerrini, R. Leonardi, A. Ortega, Symmetry-based graph Fourier transforms: are they optimal for image compression?, in: Proc. 28th IEEE Int. Conf. on Image Proc. (ICIP), 2021, pp. 1594–1598, <http://dx.doi.org/10.1109/ICIP42928.2021.9506636>.
- [5] A. Gnutti, F. Guerrini, R. Leonardi, A symmetric graph transformation framework for image compression, IEEE Access 12 (2024) 90738–90749, <http://dx.doi.org/10.1109/ACCESS.2024.3421346>.
- [6] A. Gnutti, F. Guerrini, R. Leonardi, A. Ortega, Variable-size symmetry-based graph fourier transforms for image compression, IEEE Transactions on Circuits and Systems for Video Technology 35 (12) (2025) 12772–12787, <http://dx.doi.org/10.1109/TCSVT.2025.3587753>.
- [7] J. Ballé, V. Laparra, E.P. Simoncelli, End-to-end optimized image compression, in: International Conference on Learning Representations, 2017.
- [8] J. Ballé, D. Minnen, S. Singh, S.J. Hwang, N. Johnston, Variational image compression with a scale hyperprior, in: International Conference on Learning Representations, 2018.
- [9] R. Zou, C. Song, Z. Zhang, The devil is in the details: Window-based attention for image compression, Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (2022).
- [10] Y. Zhu, Y. Yang, T. Cohen, Transformer-based transform coding, in: International Conference on Learning Representations, 2022.
- [11] F. Chen, Y. Xu, L. Wang, Two-stage octave residual network for end-to-end image compression, in: Proceedings of AAAI Conference on Artificial Intelligence, 2022.
- [12] D. He, Z. Yang, W. Peng, R. Ma, H. Qin, Y. Wang, ELIC: Efficient learned image compression with unevenly grouped space-channel contextual adaptive coding, in: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 2022.
- [13] D. Wang, W. Yang, Y. Hu, J. Liu, Neural data-dependent transform for learned image compression, in: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 2022, pp. 17379–17388.
- [14] J. Liu, H. Sun, J. Katto, Learned image compression with mixed transformer-CNN architectures, in: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR, 2023, pp. 14388–14397.
- [15] J.V.E. Team, VVC official test model vtm., 2021.
- [16] E. Alshina, J. Ascenso, T. Ebrahimi, JPEG AI: The first international standard for image coding based on an end-to-end learning-based approach, IEEE MultiMedia 31 (4) (2024) 60–69.
- [17] N. Wang, Z. Cui, Y. Su, Y. Lan, Y. Xue, C. Zhang, A. Li, Weakly supervised image dehazing via physics-based decomposition, IEEE Trans. Circuits Syst. Video Technol. 36 (1) (2026) 637–652.
- [18] Y. Su, N. Wang, Z. Cui, Y. Cai, C. He, A. Li, Real scene single image dehazing network with multi-prior guidance and domain transfer, IEEE Trans. Multimed. 27 (2025) 5492–5506.
- [19] Y.-H. Chen, K.-W. Ho, S.-R. Tsai, G.-H. Lin, A. Gnutti, W.-H. Peng, R. Leonardi, Transformer-based learned image compression for joint decoding and denoising, in: 2024 Picture Coding Symposium, PCS, 2024, pp. 1–5.
- [20] A. Gnutti, A. Alkhateeb, C.-H. Kao, E.D. Cannas, S. Mandelli, W.-H. Peng, R. Leonardi, P. Migliorati, Learned image compression for deepfake detection, in: 2025 33rd European Signal Processing Conference, EUSIPCO, 2025, pp. 1377–1381.
- [21] Y.-H. Chen, Y.-C. Weng, C.-H. Kao, C. Chien, W.-C. Chiu, W.-H. Peng, TransTIC: Transferring transformer-based image compression from human perception to machine perception, in: Proceedings of the IEEE/CVF International Conference on Computer Vision, ICCV, 2023, pp. 23297–23307.
- [22] A. Alkhateeb, A. Gnutti, F. Guerrini, R. Leonardi, J. Ascenso, F. Pereira, JPEG AI compressed domain face detection, in: 2024 IEEE 26th International Workshop on Multimedia Signal Processing, MMSP, 2024, pp. 1–6.
- [23] A. Alkhateeb, A. Gnutti, F. Guerrini, R. Leonardi, J. Ascenso, F. Pereira, JPEG AI compressed domain face detection: A multi-scale bridging perspective, IEEE Trans. Multimed. 28 (2026) 2059–2068.
- [24] C.-H. Kao, C. Chien, Y.-J. Tseng, Y.-H. Chen, A. Gnutti, S.-Y. Lo, W.-H. Peng, R. Leonardi, Bridging compressed image latents and multimodal large language models, in: The Thirteenth International Conference on Learning Representations, 2025, pp. 40042–40062.
- [25] Y. Xue, G. Jin, T. Shen, L. Tan, N. Wang, J. Gao, Y. Yu, S. Tian, Target-distractor aware UAV tracking via global agent, IEEE Trans. Intell. Transp. Syst. 26 (10) (2025) 16116–16127.

- [26] Y. Xue, G. Jin, B. Zhong, T. Shen, L. Tan, C. Xue, Y. Zheng, FMTrack: Frequency-aware interaction and multi-expert fusion for RGB-T tracking, *IEEE Trans. Circuits Syst. Video Technol.* 36 (2) (2026) 1655–1667.
- [27] Y. Baek, B. Lee, D. Han, S. Yun, H. Lee, Character region awareness for text detection, in: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2019, pp. 9365–9374.
- [28] P. He, W. Huang, T. He, Q. Zhu, Y. Qiao, X. Li, Single shot text detector with regional attention, in: *Proceedings of the IEEE International Conference on Computer Vision*, 2017, pp. 3047–3055.
- [29] X. Zhou, C. Yao, H. Wen, Y. Wang, S. Zhou, W. He, J. Liang, East: an efficient and accurate scene text detector, in: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2017, pp. 5551–5560.
- [30] M. Jaderberg, K. Simonyan, A. Zisserman, et al., Spatial transformer networks, in: *NIPS*, 2015, pp. 2017–2025.
- [31] K. He, X. Zhang, S. Ren, J. Sun, Deep residual learning for image recognition, in: *CVPR*, 2016, pp. 770–778.
- [32] B. Shi, X. Bai, C. Yao, An end-to-end trainable neural network for image-based sequence recognition and its application to scene text recognition, in: *TPAMI*, vol. 39, (11) IEEE, 2017, pp. 2298–2304.
- [33] B. Shi, X. Wang, P. Lyu, C. Yao, X. Bai, Robust scene text recognition with automatic rectification, in: *CVPR*, 2016, pp. 4168–4176.
- [34] Z. Cheng, F. Bai, Y. Xu, G. Zheng, S. Pu, S. Zhou, Focusing attention: Towards accurate text recognition in natural images, in: *ICCV*, 2017, pp. 5086–5094.
- [35] A. Graves, S. Fernández, F. Gomez, J. Schmidhuber, Connectionist temporal classification: labelling unsegmented sequence data with recurrent neural networks, in: *ICML*, 2006, pp. 369–376.
- [36] M. Lu, F. Chen, S. Pu, Z. Ma, High-efficiency lossy image coding through adaptive neighborhood information aggregation, 2022, arXiv preprint [arXiv:2204.11448](https://arxiv.org/abs/2204.11448).
- [37] J. Baek, G. Kim, J. Lee, S. Park, D. Han, S. Yun, S.J. Oh, H. Lee, What is wrong with scene text recognition model comparisons? dataset and model analysis, in: *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2019, pp. 4715–4723.
- [38] <https://github.com/Belval/TextRecognitionDataGenerator>.
- [39] V.I. Levenshtein, et al., Binary codes capable of correcting deletions, insertions, and reversals, in: *Soviet Physics Doklady*, vol. 10, (8) Soviet Union, 1966, pp. 707–710.
- [40] A. Gautam, A. Pawar, A. Joshi, S.N. Tazi, S. Chaudhary, P. Hambarde, A. Dudhane, S. Vipparthi, S. Murala, Pureformer: Transformer-based image denoising, in: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*, 2025, pp. 1441–1449.