

Character-Aware Models Improve Visual Text Rendering

Rosanne Liu*, Dan Garrette*, Chitwan Saharia, William Chan, Adam Roberts,
Sharan Narang, Irina Blok, RJ Mical, Mohammad Norouzi, Noah Constant*

Google Research

{rosanneliu, dhgarrette, nconstant}@google.com

Abstract

Current image generation models struggle to reliably produce well-formed visual text. In this paper, we investigate a key contributing factor: popular text-to-image models lack character-level input features, making it much harder to predict a word’s visual makeup as a series of glyphs. To quantify this effect, we conduct a series of experiments comparing character-aware vs. character-blind text encoders. In the text-only domain, we find that character-aware models provide large gains on a novel spelling task (WikiSpell). Applying our learnings to the visual domain, we train a suite of image generation models, and show that character-aware variants outperform their character-blind counterparts across a range of novel text rendering tasks (our DrawText benchmark). Our models set a much higher state-of-the-art on visual spelling, with 30+ point accuracy gains over competitors on rare words, despite training on far fewer examples.

1 Introduction

Over the last year, image generation models have made impressive quality gains (Rombach et al., 2021; Ramesh et al., 2022; Saharia et al., 2022; Yu et al., 2022). While many practical use cases are already within reach, *rendering visual text* in images remains a challenge. Ramesh et al. (2022) observe that DALL·E-2 “struggles at producing coherent text,” and the latest release of Stable Diffusion lists “cannot render legible text” as a known limitation.¹

In this paper, we seek to understand and improve the ability of image generation models to render high-quality visual text. To do so, we first investigate the spelling ability of text encoders in isolation. We find that despite their popularity, *character-blind* text encoders—which receive no direct signal as to the character-level makeup of their inputs—have limited spelling ability. Building on Itzhak

and Levy (2022), we test the spelling ability of text encoders across scales, architectures, input representations, languages, and tuning methods. We document for the first time the miraculous ability of character-blind models to induce robust spelling knowledge (>99% accuracy) through web pretraining, but show that this does not generalize well beyond English, and is only achieved at scales over 100B parameters, making it infeasible for most applications. We find that *character-aware* text encoders, on the other hand, are able to achieve robust spelling ability at far smaller scales.

Applying these findings to image generation, we train a range of character-aware text-to-image models and demonstrate that they significantly outperform character-blind models on text rendering. For *purely* character-level models, this improved text rendering comes at a cost—decreasing image-text alignment for prompts that don’t involve visual text. To alleviate this, we propose combining character-level and token-level input representations, and find that this delivers the best of both worlds.

Our main contributions are to: (1) Measure the spelling ability of a range of text encoders, pulling apart the effects of scale, character-awareness, and multilinguality, using a new benchmark: WikiSpell. (2) Present DrawText, the first detailed benchmark of visual text rendering for text-to-image models. (3) Improve the state of the art in text rendering ability of image generation models through the use of character-aware text encoders. We release code to reproduce our WikiSpell and DrawText evaluations.²

2 The spelling miracle

Language models can be categorized as to whether they have direct access to the characters making up their text input (“character-aware”) or do not (“character-blind”). Many early neural lan-

*Equal contribution.

¹<https://huggingface.co/stabilityai/stable-diffusion-2-1>

²<https://github.com/google-research/google-research/tree/master/drawtext>



Figure 1: **Top:** Image generation models lacking character-level input features often misspell words. **Bottom:** Using a character-aware text encoder significantly improves the accuracy of rendered text. Prompts are: *A vintage postage stamp with the message: _____, with messages: (1) California: All Dreams Welcome, (2) Canada: For Glowing Hearts, (3) Colorado: It's Our Nature, (4) St. Louis: All Within Reach.*

language models operated directly on characters, with no notion of multi-character “tokens” (Sutskever et al., 2011; Graves, 2013). Later models moved to vocabulary-based tokenization, with some like ELMo (Peters et al., 2018) retaining character-awareness, and others like BERT (Devlin et al., 2019) abandoning it in favor of more efficient pretraining. At present, most widely used language models are character-blind, relying on data-driven subword segmentation algorithms like Byte Pair Encoding (BPE) (Gage, 1994; Sennrich et al., 2016) to induce a vocabulary of subword pieces. While these methods back off to character-level representations for sufficiently *rare* sequences, they compress *common* character sequences into unbreakable units by design, as shown in Figure 2.

Recent work on “token-free” modeling has pointed to advantages of character-aware input representations. Xue et al. (2022) show that ByT5—a character-aware multilingual language model trained directly on UTF-8 bytes—outperforms parameter-matched character-blind models on tasks related to spelling and pronunciation. While operating at the byte or character level comes at the cost of training and inference speed, additional work suggests that this can be overcome through downsampling (Clark et al., 2022; Tay et al., 2021). See Mielke et al. (2021) for a recent overview of tokenization methods and character awareness.

Surprisingly, despite lacking direct access to a token’s spelling, character-blind models are, to varying degree, able to *infer* the character-level




Figure 2: Subword tokenization—used in most text-to-image models—maps common character sequences onto IDs that are looked up in an embedding table before being passed to the model, removing any signal about token-internal composition. This example uses the T5 SentencePiece tokenizer (Raffel et al., 2020).

makeup of their tokens. Itzhak and Levy (2022) observe that, after fine-tuning for spelling, RoBERTa and GPT-2 can achieve 32% and 33% accuracy at spelling held-out tokens. Kaushal and Mahowald (2022) confirm this ability and probe it further; however it remains unclear where in pretraining this knowledge is coming from, and how to improve it. For example, should we expect larger character-blind models to reach 100% spelling accuracy across all tokens in their vocabulary?

In section §3 we find that, with sufficient scale, character-blind models can achieve near-perfect spelling accuracy. We dub this phenomenon the “spelling miracle”, to emphasize the difficulty of inferring a token’s spelling from its distribution alone. At the same time, we observe that character-blind text encoders of the sizes used in practice for image generation are lacking core spelling knowledge.

With this in mind, it is unsurprising that to-

day’s image generation models struggle to translate input tokens into glyph sequences. These models’ text encoders are all character-blind, with Stable Diffusion, DALL·E, DALL·E-2, Imagen, Parti and eDiff-I all adopting BPE tokenizers (Rombach et al., 2021; Ramesh et al., 2021, 2022; Saharia et al., 2022; Yu et al., 2022; Balaji et al., 2022).

For image-text models, another key source of knowledge is supervised image-caption data. Even if its text encoder is character-blind, could a model learn to spell by observing the makeup of words within images? While possible, we suspect this is an inefficient paradigm for learning, as each token would need to be learned separately, and would need to appear within an image-caption pair seen in training. In section §5 we find that, indeed, this “late-stage” learning of spelling is inferior to using a pretrained character-aware text encoder.

3 Measuring text encoder spelling ability

Since text-to-image generation models rely on text encoders to produce the representations for decoding, we first explore the ability of text encoders in isolation, using a text-only spelling evaluation task.

3.1 The WikiSpell benchmark

We build the WikiSpell benchmark by sampling words from Wiktionary.³ For each example in the dataset, the input to the model is a single word, and the expected output is its spelling, generated by inserting spaces between each Unicode character:⁴

elephant → e l e p h a n t

To examine the relationship between a word’s frequency and a model’s ability to spell it, we group words into buckets based on their frequency in the mC4 corpus (Xue et al., 2021). We create test and development sets from each bucket by sampling 1,000 words uniformly. The five buckets used are: the top 1% most common words, the 1–10% most common, 10–20%, 20–30%, and the bottom 50% (which includes words never seen in the corpus). Finally, we build a training set of 10,000 words by combining 5,000 words sampled uniformly from the bottom 50% bucket with 5,000 sampled proportional to their frequencies in mC4. We exclude words in the dev and test sets from the training set, so evaluation is always on *held out* words.

Beyond English, we repeat this process for six other languages: Arabic, Chinese, Finnish, Korean,

Russian, and Thai. For language selection criteria and further technical details, see Appendix D.

The WikiSpell benchmark is similar to SpellingBee, introduced by Itzhak and Levy (2022), but differs in a few key ways. First, SpellingBee is designed to probe a model’s *embedding matrix*: given an embedding vector, SpellingBee seeks to output the character sequence of the corresponding vocabulary element. As such, SpellingBee’s inputs are always a single token, and it does not measure spelling ability for words the model represents as multiple tokens. Second, due to how subword vocabularies are trained, model vocabularies only contain high-frequency words, and thus SpellingBee inputs are necessarily high-frequency. Finally, as inputs must be drawn from a model’s vocabulary, SpellingBee training and evaluation data must be tailored to a specific model, so a dataset cannot be reused across models. In contrast, WikiSpell is model-agnostic, covers single- to many-token words, and covers high- to low-frequency words.

3.2 Text generation experiments

We use the WikiSpell benchmark to evaluate pretrained text-only models across a variety of scales. In particular, we experiment with: **T5** (Raffel et al., 2020), a character-blind encoder-decoder model pretrained on English data; **mT5** (Xue et al., 2021), which is similar to T5, but pretrained on >100 languages; **ByT5** (Xue et al., 2022), a character-aware version of mT5 that operates directly on UTF-8 byte sequences; and **PaLM** (Chowdhery et al., 2022), a decoder-only model of much larger scale, pretrained predominantly on English. Experimental results from English-only evaluation are shown in Table 1, and multilingual evaluation in Table 2.

Our first finding is that character-blind models T5 and mT5 perform much worse on the Top-1% most frequent words. This result may seem counter-intuitive since models typically perform *better* on frequent items; however due to how subword vocabularies are trained, common words are typically represented as atomic tokens (e.g., 87% of words in the English Top 1% bucket are single-token for T5), thereby obscuring their internal makeup. Scores are a bit higher in the mid-frequency buckets, where words are typically broken into a few common tokens, and lower again in the lowest-frequency bucket, where even the subword tokens may be less frequent. The low spelling accuracies indicate that T5’s encoder does not retain sufficient information

³<https://www.wiktionary.org/>

⁴i.e., in Python 3: " ".join(word)

Frequency Bucket	Fine-tuned, with frozen encoder								Fine-tuned, all parameters trained								Few-shot		
	T5				ByT5				T5				ByT5				PaLM		
	B	L	XL	XXL	B	L	XL	XXL	B	L	XL	XXL	B	L	XL	XXL	8B	62B	540B
Top 1%	14	12	50	66	97	95	97	98	36	46	62	68	99	100	100	100	84	99	100
1–10%	29	24	67	69	97	95	98	98	67	72	82	85	100	100	100	100	62	98	99
10–20%	35	27	73	73	96	94	98	98	74	79	89	91	100	100	100	100	70	97	99
20–30%	32	24	68	68	96	94	99	98	74	78	87	90	100	100	100	100	71	97	99
Bottom 50%	29	22	64	65	97	95	99	98	75	77	88	90	100	100	100	100	69	97	99

Table 1: WikiSpell exact-match accuracy results for English. T5 models range from Base (B) (250M params) to XXL (11B params), while ByT5 models range from Base (300M) to XXL (13B).

Language	Fine-tuned, with frozen encoder								Few-shot		
	mT5				ByT5				PaLM		
	B	L	XL	XXL	B	L	XL	XXL	8B	62B	540B
Arabic	22	60	75	87	99	99	100	99	32	68	89
Chinese	78	76	83	84	99	98	99	99	81	93	98
English	7	32	54	71	98	96	99	99	71	97	99
Finnish	10	36	62	77	98	97	99	99	45	84	99
Korean	37	58	77	81	99	99	100	99	71	88	96
Russian	9	41	57	76	99	98	99	99	41	86	98
Thai	29	42	46	60	99	99	99	99	22	39	63
Average	27	49	65	77	99	98	99	99	52	79	92

Table 2: WikiSpell exact-match accuracy results on 7 diverse languages, averaged across all frequency buckets. mT5 and ByT5 models were fine-tuned on the combined training sets of all languages; PaLM was prompted with 20 in-language examples.

about the spelling of subwords in its vocabulary.

Secondly, we find that for character-blind models, scale is a key factor in spelling ability. Both T5 and mT5 improve with scale, but even at XXL size, they are not particularly strong (e.g., T5-XXL’s performance on common English words is only 66%). Only when character-blind models reach PaLM’s scale do we start to see near-perfect spelling ability: PaLM 540B achieves >99% accuracy across all frequency buckets in English, despite the fact that it sees only 20 examples in its prompt (as opposed to the 1,000 fine-tuning examples shown to T5). However, performance is lower on other languages.

Our experiments on ByT5 show that character-aware models have far greater spelling ability. ByT5’s performance at Base and Large sizes is only slightly behind XL and XXL (though still in at least the mid-90% range), and the frequency of a word has little effect on ByT5’s ability to spell it. These results far exceed those of (m)T5, and are comparable to the English performance of PaLM, which has >100× more parameters, and exceed PaLM’s performance on other languages. These findings indicate that substantially more character-level information is retained by the ByT5 encoder, and in such a way that it can be retrieved from those

frozen parameters as needed for the decoding task.

We also test fine-tuning the full model instead of keeping the encoder frozen (also in Table 1). When ByT5’s encoder is fine-tuned, performance goes to roughly 100% across all scales and frequency buckets. For T5, the effect of full fine-tuning is more mixed: on rare words, it helps a lot (65% → 90% for T5-XXL on Bottom 50%), while for common words, it has little effect (66% → 68% on Top 1%). This tells us that for words that get broken into smaller pieces (which may be repeated across training examples), the model can memorize spelling information provided during fine-tuning, whereas for single-token words, fine-tuning provides no spelling signal since, by definition, that single token will not appear in the fine-tuning dataset.

4 The DrawText benchmark

Evaluating text-to-image models has been an ongoing topic of research, with the development of standard benchmarks from COCO (Lin et al., 2014) to DrawBench (Saharia et al., 2022), and metrics including FID (Heusel et al., 2017), CLIP score (Heusel et al., 2021), and human preferences (Saharia et al., 2022). However, there has been a lack of work on text rendering and spelling evaluation. To that end, we present a new benchmark, DrawText, designed to measure the text rendering quality of text-to-image models. The benchmark consists of two parts, assessing distinct model capabilities: 1) DrawText Spelling, which evaluates simple word rendering over a large set of English terms; and 2) DrawText Creative, which evaluates end-to-end text rendering across diverse visual settings.

4.1 DrawText Spelling

To measure the spelling ability of image generation models in a controlled and automatable fashion, we construct 500 prompts by sampling 100 words from each of the English WikiSpell frequency buckets

(see §3.1), and plugging them into the template: *A sign with the word “_____” written on it.* For each prompt, we sample 4 images from the candidate model, and assess them using optical character recognition (OCR)-based metrics.⁵

For OCR evaluation, we use the Google Cloud Vision API, which returns all text found within an image, along with bounding box locations. The DrawText Spelling prompt tends to generate a prominently positioned sign with text, which is relatively simple for off-the-shelf OCR to identify, but if the system returns multiple bounding boxes, we only use the top-most one. Additionally, since text may be rendered across multiple lines, we post-process the OCR output by removing newline characters that appear within a single bounding box. Finally, since text on real signs is often written in all capitals, and models often do the same regardless of how the word is written in the prompt, we ignore case when computing the spelling accuracy.

4.2 DrawText Creative

Visual text is not limited to mundane examples like street signs. Text can appear in many forms—scribbled, painted, carved, sculpted, and so on. If image generation models support flexible and accurate text rendering, this can help designers in developing creative fonts, logos, layouts, and more.

To test the ability of image generation models to support these use cases, we worked with a professional graphic designer to construct 175 diverse prompts that require rendering text in a range of creative styles and settings. The prompts vary in how much text is specified, ranging from a single letter to an entire sentence. We share these prompts in Appendix G, with the expectation that they will help the community work towards improving text rendering. Many of the prompts are beyond the abilities of current models, with state-of-the-art models exhibiting misspelled, dropped, or repeated words, as seen in Figure 3.

5 Image generation experiments

In this section, we evaluate the spelling ability of text-to-image generative models with the proposed DrawText benchmark. State-of-the-art text-to-image generative models consist of a text encoder plus a cascade of either diffusion models (Sa-

⁵For models supporting non-square inference, we use 2:1 aspect ratio, for better visualization. This gave a modest gain in OCR performance in preliminary experiments. We rescale all images to 64×64 before running OCR for fair comparison.

haria et al., 2022) or autoregressive models (Yu et al., 2022) that map the encoded text representations to realistic images. In section §3 we saw that character-aware text encoders greatly outperform character-blind models on spelling in a text-only setting; in this section, we investigate whether making the text encoder character-aware improves the text rendering ability of text-to-image models.

5.1 Models

For an apples-to-apples comparison, we train two character-blind and three character-aware image generation models. Our training closely follows the procedure of Saharia et al. (2022), with the following modifications. First, our models train for 500,000 steps, which is $5.6 \times$ fewer steps than Imagen. Second, we only train the initial 64×64 model, as text rendering ability can already be assessed at this scale. This allows us to forgo the training of super-resolution models.

Third, rather than a mixture of datasets, we train exclusively on the publicly available Laion-400M (Schuhmann et al., 2021). This improves reproducibility and also increases the amount of visual text seen during training. Inspecting a random sample of 100 images, we found that a relatively high proportion (around 71%) of Laion images contain text, and many (around 60%) exhibit correspondence between caption text and visual text.

Fourth, to prevent models from clipping text, we train on *uncropped* images with arbitrary aspect ratios. In contrast with the widely used strategy of cropping a square from the center of the image, we maintain the image’s true aspect ratio by padding with black borders. The model receives an additional binary mask input indicating the padding.⁶

To test the effects of text encoder size and character-awareness, we vary the pretrained text encoder as follows:

T5-XL and T5-XXL — Following Saharia et al. (2022), we use the (character-blind) pretrained T5 text encoders of Raffel et al. (2020). The encoder sizes are 1.2B (XL) and 4.6B (XXL). Note, T5-XXL is the same encoder used in both Imagen and the recent eDiff-I (Balaji et al., 2022).

ByT5-XL and ByT5-XXL — We use the pretrained ByT5 encoders of Xue et al. (2022), with encoders sizes 2.6B (XL) and 9.0B (XXL). These differ from T5 in several regards. First, ByT5

⁶We apply the above strategy for 80% of training examples, and use center cropping for the remaining 20%.



Figure 3: Sample generations from Imagen (Saharia et al., 2022) on eight DrawText Creative prompts, illustrating their diversity. Most samples exhibit misspelled text or misshapen glyphs (e.g., the ff in coffee). Errors are increasingly common with prompts seeking longer text spans. See Appendix G for all prompts.

models read and write UTF-8 bytes rather than tokens from a vocabulary, so they are fully character-aware. Second, ByT5 is multilingual, trained on the mC4 corpus of over 100 languages. Third, ByT5 pretrains with sequence length 1024, twice that of T5. When encoding text as input to the image generation module, we use a sequence length of 256 bytes, compared to 64 tokens for the T5 models.

Concat(T5-XXL, ByT5-Small) — We use as the text encoding a concatenation of the encodings from T5-XXL and a small ByT5 model. ByT5-Small (220M) represents a lightweight addition to the Saharia et al. (2022) model in terms of overall compute and model size (a 4.8% increase in encoder size), but makes the model character-aware.

Imagen Aspect-Ratio (Imagen-AR) — To test the benefit of training on uncropped images, we fine-tune the Imagen model of Saharia et al. (2022) for an additional 380,000 steps, to 3.2M steps total, training on uncropped images with preserved original aspect ratio, as described above.

Beyond these custom models, we benchmark Stable Diffusion version 1.5 (Rombach et al., 2021), DALL-E-2 (Ramesh et al., 2022), Parti (Yu et al., 2022) and Imagen (Saharia et al., 2022), all of which use character-blind text encoders. Among these, Imagen is most similar to our experimental models, using the same T5-XXL encoder, but trained much longer and with a larger scale of data.

5.2 DrawText Spelling results

Char-aware models improve spelling. Figure 4 shows our DrawText Spelling results across 10 models, with 2,000 images sampled per model, evaluated with OCR. Accuracy is computed on the full string (i.e., no credit is given for partial matches). Across all word frequencies, character-aware models (ByT5 and Concat) outperform the rest, with 15+ point accuracy gains over Imagen-AR on the most frequent words, and 30+ point gains on the least frequent words. This is remarkable given that Imagen-AR trained for $6.6\times$ longer.

Our T5 models (character-blind) provide a more controlled comparison against the character-aware models, as they differ *only* in the choice of text encoder—trained on the same dataset for the same number of steps. Here, the gains are even larger: 25+ point gains on the most frequent words and 30+ point gains on the least frequent. Notably, these gains persist even for the *smaller* ByT5-XL model, whose encoder is 43% smaller than T5-XXL.

To estimate the rate of OCR errors, we manually validate a balanced set of 128 samples from T5-XXL and ByT5-XXL (see Appendix B for details). We find no false positives, but when OCR detects an error, ByT5-XXL is actually correct 34% of the time, while T5-XXL is correct 9%. This asymmetry suggests the benefit of character-aware modeling may be even greater than implied by Figure 4.




Figure 4: Accuracy of 10 image generation models on DrawText Spelling. Character-aware models (ByT5 and Concat) outperform others regardless of size, and particularly on rare words. Imagen-AR shows the benefit of avoiding cropping, but still underperforms character-aware models, despite training $6.6\times$ longer.




Error Type	Examples
Semantic ✓	demonstrated → demonstrafied inquisitiveness → inquisitioness
Homophone ✓	accommodate → accomidate Toronto → Torondo
Add Glyph ✓	labor → labort debut → debust
Drop Glyph ✗	stopping → stoping experiments → experimets
Repeat Glyph ✗	possible → possible locate → locaate
Merge Glyphs ✗	
Misshape ✗	
No Text ✗	

Table 3: Error types observed on DrawText Spelling. Some errors (✓) are not observed in character-aware models. Others (✗) are found across all model types.

Char-aware models make fewer types of error.

To gain a better understanding of different models’ failure modes, we manually inspect our T5 and ByT5 model outputs. Table 3 illustrates common error types. Several categories of error are *only* observed in T5 models, suggesting that they stem from the encoder’s lack of core spelling knowledge. In **semantic** errors, the model makes a plausible morpheme substitution, as in *demonstrated* → *demonstrafied*. In **homophone** errors, the model produces an incorrect spelling that could be pronounced similarly to the target word. This suggests that some of T5’s “miraculous” spelling ability may

derive from online pronunciation guides. In **add glyph** errors, the model inserts a letter that was absent from the target, again reflecting the model’s uncertainty about a token’s internal character makeup.

One notable sub-type of semantic error is character-blind models “regularizing” irregular inflection, as in *fought* → *fighted*. On a hand-chosen set of 23 common irregular past-tense verbs (*began, chose, dug, etc.*), we find T5-based models erroneously add *-ed* in 11% of samples (see Figure 6), while our character-aware models never exhibit this type of error. This is clear evidence that character-blind models partly rely on a word’s meaning (*fought* ⇒ PAST) and fallible patterns of morphology (PAST ⇒ *-ed*) to predict spelling.

Other error categories are found across all model types; these include dropped, repeated, merged, or misshapen glyphs. Given that our ByT5 encoders provide a robust spelling signal (see §3.2), we understand these errors to be “layout issues”, where the image generation module has trouble shaping and positioning realistic glyphs within the image.

Char-aware models reduce consistent errors.

Another stark difference between our models lies in whether they *consistently* misspell a given word across multiple samples. As illustrated in Figure 5, there are many words that our T5 models misspell no matter how many samples are drawn. Again, we believe this indicates missing knowledge in the text encoder. By contrast, our ByT5 models are more likely to make sporadic errors. We quantify this observation in Figure 7 by measuring the rates at which the model is consistently right (4/4) or wrong (0/4) across all four image samples. On common words in particular (Top 1%), we see a sharp contrast in that ByT5 models are *never* consistently wrong, while T5 models are consistently




Figure 5: Selected errors from two models on DrawText Spelling. **Top:** Sample words that T5-XXL consistently misspells (target: *refers*, *similarly*, *stomach*). **Bottom:** ByT5-XXL errors (red outline) are more sporadic and minor: dropped, merged, or repeated glyphs. See Appendix F for more examples.



Figure 6: Character-blind T5 models erroneously use -(e)d endings for irregular past tense verbs (targets: *fought*, *threw*, *knew*, *chose*).




Figure 7: Proportion of words for which models are consistently correct (top) or consistently incorrect (bottom) across 4 image samples on DrawText Spelling.




Figure 8: On DrawText Creative, our Concat model is preferred over T5-XXL on fidelity and alignment, as well as a new metric assessing rendered text accuracy.

wrong on 10% or more of words.

5.3 DrawText Creative results

To test our models in a more realistic user-facing setting, we sample 8 images from each of our T5 and ByT5 models on our 175 DrawText Creative prompts in Appendix G. These prompts are more diverse and challenging, with the majority targeting three or more words of rendered text.

Focusing on text rendering ability,⁷ we find once again that character-aware models have a clear advantage. Figures 12 and 13 show non-cherrypicked samples on two prompts where T5-XXL consistently misspells one or more words. On prompts targeting longer text spans, all our models struggle, as seen in Figure 14. Nevertheless, we observe that character-aware text encoders provide a clear lift on these prompts, reducing the misspellings of words like *refrain*, *arguing*, and *chimpanzees*.

We confirm the above observations quantitatively by comparing T5-XXL vs. Concat using the DrawBench methodology (Saharia et al., 2022), evaluated over our 175 creative prompts. Beyond the standard DrawBench metrics of fidelity and alignment (described in the following section), we ask raters *Which set of images more accurately shows the text: “<target text>”?* Results in Figure 8 show Concat is preferred on all metrics.

5.4 DrawBench results

We have shown that character-aware text encoders excel at spelling, in both text (§3) and visual (§5) domains. But does this ability come at a cost? Can these models maintain competitive image quality and text-image alignment, even on prompts that

⁷We note our models’ overall image quality and alignment fall short of a state-of-art model like Imagen (see Figure 3). This is expected, given that our models train exclusively on the lightly curated Laion-400M dataset (Schuhmann et al., 2021), and see $5.6\times$ fewer examples than Imagen during training.




Figure 9: DrawBench user preference rates comparing T5-XXL with three character-aware models. While image fidelity is comparable, the pure ByT5 models have lower image-text alignment. Concat closes the alignment gap. Error bars show 95% confidence intervals.

don’t require text rendering? To shed light on this question, we run several side-by-side comparisons using the DrawBench evaluation of Saharia et al. (2022). This asks human raters to compare two models’ generations of 8 images each across 200 prompts covering 11 thematic categories. We follow the procedure described in Saharia et al. (2022) closely, aggregating scores across 25 raters.

Figure 9 shows DrawBench results of three side-by-side comparisons of character-aware models vs. T5-XXL. While image quality (“fidelity”) is similar across the board, we find that *purely* character-level models (ByT5) score worse on image-text alignment, with raters preferring T5-XXL on 60% of prompts. By contrast, our Concat model closes this gap to within error bars. Thus, this “hybrid” character-aware model is able to greatly improve text rendering (Figure 4), without significantly hurting performance elsewhere.

Appendix C provides a per-category breakdown and further analysis; while the character-aware models excel on the DrawBench text category, the ByT5 models are dispreferred in most other categories. Through manual inspection, we find the ByT5 models are more prone to ignore information in the prompt, for example leaving out a mentioned object, or choosing a canonical color over a requested one. One possible explanation for this behavior is that we did not tune the *guidance weight* parameter used at inference time (Saharia et al., 2022), using a fixed value of 30 throughout. Increasing this parameter is known to boost image-text alignment, but at the cost of diversity. It may be that character-level models benefit from higher guidance values than token-based models.

Another possibility is that the ByT5 models have a shallower understanding of English language due to their multilingual nature—as ByT5 was exposed to roughly 70× less English than T5 during pre-

training.⁸ Given this difference, we should also expect to see corresponding gains on non-English languages. We confirm this expectation through preliminary results in Appendix A.

6 Conclusion

In this paper, we set out to better understand what is needed for image generation models to reliably render well-formed visual text. Using our novel WikiSpell and DrawText benchmarks, we were able to precisely quantify the effects of character-awareness and other design choices on spelling ability in both the text and visual domains.

We found that character-aware text encoders confer large gains on spelling, and when used within an image generation model, these gains translate into improved visual text rendering. However, using *exclusively* character-level representations deteriorated overall text-image alignment—at least when evaluating our multilingual ByT5 text encoder on English prompts with untuned guidance weight. To resolve this, we found that a hybrid model combining token-level and character-level signals offered the best of both worlds: dramatically improving visual text without significantly affecting alignment.

While we saw substantial improvements on DrawText Spelling accuracy (75% → 94% on common words and 47% → 83% on rare words), some failure modes remain unaddressed. Even our strongest models were observed to occasionally drop, repeat, or merge letters within a word, or words within a phrase. Our results strongly suggest that resolving these issues will require orthogonal improvements outside the text encoder, specifically changes to the image generation module.

As a secondary finding, we demonstrated for the first time that, with sufficient scale, even models lacking a direct character-level view of their inputs can *infer* robust spelling information through knowledge gained via web pretraining—“the spelling miracle”. While remarkable, this finding is less immediately practical, as it requires models over 100B parameters, and even these did not generalize well beyond English in our experiments.

Limitations

While we establish the “miraculous” ability of character-blind models to induce robust spelling

⁸The models were trained on the same number of tokens, but only 6% of ByT5 training was on English, and we estimate 4 UTF-8 bytes per T5 token.

information through large-scale web pretraining, our work does not attempt to identify the mechanisms or sources through which this information is learned. Possible sources within web corpora include: dictionaries containing phonetic pronunciation guides, alphabetically ordered lists, typos and other misspellings, and examples of spelling words with dashes or spaces between every character. Linguistic phenomena that may aide in inducing spelling knowledge include words with predictable morphemic makeup, and cases where meaning-form relation is non-arbitrary, contra Saussure’s “semiotic arbitrariness”. We refer the reader to [Itzhak and Levy \(2022\)](#) and [Kaushal and Mahowald \(2022\)](#) for work in this direction.

Most of our image generation experiments are limited to English. We present preliminary results in [Appendix A](#) showing that our ByT5-based models have stronger multilingual understanding than T5. However it would be valuable to test this further, and to explore training image generation models on multilingual image-caption datasets, as opposed to merely using a pretrained multilingual text encoder.

Ideally, it would be possible to conduct controlled comparisons between pretrained text encoders that differ only in one regard, to isolate all factors contributing to performance. However as pretraining large language models is resource intensive, we were only able to use off-the-shelf text encoders, which often differ along multiple axes. In our text-only experiments, we isolated the contributions of character-awareness (ByT5 vs. mT5/T5) and multilinguality (ByT5/mT5 vs. T5). However, in our image generation experiments, these factors were conflated, as we had limited resources for training new models. Still, the fact that ByT5-based image generation models outperform T5 *despite* being multilingual (which often degrades performance on English-only tasks) strongly suggests that character-awareness is the key factor for spelling ability.

Another limitation is that we focused on image generation models that leverage *frozen* pretrained text encoders. This enabled straightforward experimentation by swapping encoders and retraining the image generation module. However, it remains to be seen whether our results extend to settings where the text encoder is trained along with the rest of the model, as in [Yu et al. \(2022\)](#).

Ethics Statement

We note our image-caption training data comes from Laion-400M ([Schuhmann et al., 2021](#)), which is uncurated and known to contain harmful biases and offensive content. We hope to utilize and contribute to safer and better curated datasets in the future, as well as to develop improved techniques for debiasing and detoxifying existing models.

A potential risk for image generation models is that they can be used for creating misleading and harmful content. Our work on improving text rendering could aide the creation of fake signs and other misleading images containing visual text. With the wide-spread availability of image generation models, we expect that improving education around misinformation and adopting better digital signature mechanisms will be important countermeasures.

Acknowledgements

We thank Jason Baldrige, Jon Clark, Noah Fiedel, Linting Xue, and Jiahui Yu for helpful discussion and comments on an earlier draft. We thank Sarah Pratt for validating findings on Stable Diffusion models.

References

- Yogesh Balaji, Seungjun Nah, Xun Huang, Arash Vahdat, Jiaming Song, Karsten Kreis, Miika Aittala, Timo Aila, Samuli Laine, Bryan Catanzaro, Tero Karras, and Ming-Yu Liu. 2022. [eDiff-I: Text-to-image diffusion models with an ensemble of expert denoisers](#). *CoRR*, abs/2211.01324.
- Aakanksha Chowdhery, Sharan Narang, Jacob Devlin, Maarten Bosma, Gaurav Mishra, Adam Roberts, Paul Barham, Hyung Won Chung, Charles Sutton, Sebastian Gehrmann, Parker Schuh, Kensen Shi, Sasha Tsvyashchenko, Joshua Maynez, Abhishek Rao, Parker Barnes, Yi Tay, Noam Shazeer, Vinodkumar Prabhakaran, Emily Reif, Nan Du, Ben Hutchinson, Reiner Pope, James Bradbury, Jacob Austin, Michael Isard, Guy Gur-Ari, Pengcheng Yin, Toju Duke, Anselm Levskaya, Sanjay Ghemawat, Sunipa Dev, Henryk Michalewski, Xavier Garcia, Vedant Misra, Kevin Robinson, Liam Fedus, Denny Zhou, Daphne Ippolito, David Luan, Hyeontaek Lim, Barret Zoph, Alexander Spiridonov, Ryan Sepassi, David Dohan, Shivani Agrawal, Mark Omernick, Andrew M. Dai, Thanumalayan Sankaranarayanan Pillai, Marie Pellat, Aitor Lewkowycz, Erica Moreira, Rewon Child, Oleksandr Polozov, Katherine Lee, Zongwei Zhou, Xuezhi Wang, Brennan Saeta, Mark Diaz, Orhan Firat, Michele Catasta, Jason Wei, Kathy Meier-Hellstern, Douglas Eck,

- Jeff Dean, Slav Petrov, and Noah Fiedel. 2022. [Palm: Scaling language modeling with pathways](#). *arxiv:2204.02311*.
- Jonathan H. Clark, Dan Garrette, Iulia Turc, and John Wieting. 2022. [Canine: Pre-training an efficient tokenization-free encoder for language representation](#). *Transactions of the Association for Computational Linguistics*, 10:73–91.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [BERT: Pre-training of deep bidirectional transformers for language understanding](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.
- Philip Gage. 1994. A new algorithm for data compression. *C Users J.*, 12(2):23–38.
- Alex Graves. 2013. [Generating sequences with recurrent neural networks](#). *CoRR*, abs/1308.0850v5.
- Jack Hessel, Ari Holtzman, Maxwell Forbes, Roman Le Bras, and Yejin Choi. 2021. [CLIPScore: A reference-free evaluation metric for image captioning](#). *CoRR*, abs/2104.08718.
- Martin Heusel, Hubert Ramsauer, Thomas Unterthiner, Bernhard Nessler, and Sepp Hochreiter. 2017. GANs trained by a two time-scale update rule converge to a local Nash equilibrium. *Advances in neural information processing systems*, 30.
- Itay Itzhak and Omer Levy. 2022. [Models in a spelling bee: Language models implicitly learn the character composition of tokens](#). In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 5061–5068, Seattle, United States. Association for Computational Linguistics.
- Ayush Kaushal and Kyle Mahowald. 2022. [What do tokens know about their characters and how do they know it?](#) In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 2487–2507, Seattle, United States. Association for Computational Linguistics.
- Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. 2014. Microsoft COCO: Common objects in context. In *European conference on computer vision*, pages 740–755. Springer.
- Sabrina J. Mielke, Zaid Alyafeai, Elizabeth Salesky, Colin Raffel, Manan Dey, Matthias Gallé, Arun Raja, Chenglei Si, Wilson Y. Lee, Benoît Sagot, and Samson Tan. 2021. [Between words and characters: A brief history of open-vocabulary modeling and tokenization in NLP](#). *CoRR*, abs/2112.10508.
- Matthew E. Peters, Mark Neumann, Mohit Iyyer, Matt Gardner, Christopher Clark, Kenton Lee, and Luke Zettlemoyer. 2018. [Deep contextualized word representations](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 2227–2237, New Orleans, Louisiana. Association for Computational Linguistics.
- Telmo Pires, Eva Schlinger, and Dan Garrette. 2019. [How multilingual is multilingual BERT?](#) In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 4996–5001, Florence, Italy. Association for Computational Linguistics.
- Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J. Liu. 2020. [Exploring the limits of transfer learning with a unified text-to-text transformer](#). *Journal of Machine Learning Research*, 21(140):1–67.
- Aditya Ramesh, Prafulla Dhariwal, Alex Nichol, Casey Chu, and Mark Chen. 2022. Hierarchical text-conditional image generation with CLIP latents. *arXiv preprint arXiv:2204.06125*.
- Aditya Ramesh, Mikhail Pavlov, Gabriel Goh, Scott Gray, Chelsea Voss, Alec Radford, Mark Chen, and Ilya Sutskever. 2021. [Zero-shot text-to-image generation](#). In *Proceedings of the 38th International Conference on Machine Learning*, volume 139 of *Proceedings of Machine Learning Research*, pages 8821–8831. PMLR.
- Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. 2021. [High-resolution image synthesis with latent diffusion models](#). *CoRR*, abs/2112.10752.
- Chitwan Saharia, William Chan, Saurabh Saxena, Lala Li, Jay Whang, Emily Denton, Seyed Kamyar Seyed Ghasemipour, Burcu Karagol Ayan, S Sara Mahdavi, Rapha Gontijo Lopes, et al. 2022. Photorealistic text-to-image diffusion models with deep language understanding. *arXiv preprint arXiv:2205.11487*.
- Christoph Schuhmann, Richard Vencu, Romain Beaumont, Robert Kaczmarczyk, Clayton Mullis, Aarush Katta, Theo Coombes, Jenia Jitsev, and Aran Komatsuzaki. 2021. [LAION-400M: open dataset of clip-filtered 400 million image-text pairs](#). *CoRR*, abs/2111.02114.
- Rico Sennrich, Barry Haddow, and Alexandra Birch. 2016. [Neural machine translation of rare words with subword units](#). In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1715–1725, Berlin, Germany. Association for Computational Linguistics.

Ilya Sutskever, James Martens, and Geoffrey Hinton. 2011. Generating text with recurrent neural networks. In *Proceedings of the 28th International Conference on International Conference on Machine Learning, ICML'11*, page 1017–1024, Madison, WI, USA. Omnipress.

Yi Tay, Vinh Q. Tran, Sebastian Ruder, Jai Prakash Gupta, Hyung Won Chung, Dara Bahri, Zhen Qin, Simon Baumgartner, Cong Yu, and Donald Metzler. 2021. *Charformer: Fast character transformers via gradient-based subword tokenization*. *CoRR*, abs/2106.12672.

Linting Xue, Aditya Barua, Noah Constant, Rami Al-Rfou, Sharan Narang, Mihir Kale, Adam Roberts, and Colin Raffel. 2022. *ByT5: Towards a token-free future with pre-trained byte-to-byte models*. *Transactions of the Association for Computational Linguistics*, 10:291–306.

Linting Xue, Noah Constant, Adam Roberts, Mihir Kale, Rami Al-Rfou, Aditya Siddhant, Aditya Barua, and Colin Raffel. 2021. *mT5: A massively multilingual pre-trained text-to-text transformer*. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 483–498, Online. Association for Computational Linguistics.

Jiahui Yu, Yuanzhong Xu, Jing Yu Koh, Thang Luong, Gunjan Baid, Zirui Wang, Vijay Vasudevan, Alexander Ku, Yinfei Yang, Burcu Karagol Ayan, et al. 2022. Scaling autoregressive models for content-rich text-to-image generation. *arXiv preprint arXiv:2206.10789*.

A Multilingual results

As ByT5 is a multilingual model covering 100+ languages, we are interested to see if image generation models built on ByT5 deliver improved performance over T5 on non-English languages. While the text encoder itself is multilingual, it is not obvious whether this is sufficient to produce a multilingual image generation model. The image caption dataset used for training in all of our experiments is Laion-400M (Schuhmann et al., 2021), which we estimate through language ID detection to consist of 95% English captions, with only minimal coverage (<0.1%) of some widely spoken languages, such as Arabic and Hindi.

To test for multilingual understanding, we translate two English prompts to 11 languages using Google Translate, and feed the outputs to our models. As can be seen in the rows 1–2 of Figure 10, our T5-XXL model demonstrates basic understanding of five high-resource European languages (German, French, Spanish, Portuguese,

Russian).⁹ However, in the remaining languages (Greek, Hindi, Arabic, Chinese, Japanese, Korean), T5 appears to ignore the caption completely.

By comparison, our ByT5-XXL model exhibits understanding across all 11 languages. Given its limited training on multilingual captions, we interpret this ability as due to the pretrained ByT5 encoder’s alignment of representations across languages. If the encoder already embeds similar prompts into a shared space that factors out the contribution of language, then the image generation model should be able to learn from just a handful of examples how to map any language seen in pretraining into the space of images.¹⁰

If this explanation is correct, it also suggests that rendering text in different scripts will require more than just a multilingual encoder. To learn the glyph shapes, variants and fonts used for a given script, we should expect to need to train models on a large source of visual text in that script. Indeed, rows 3 and 6 of Figure 10 show that neither of our models can map prompt text onto visual text in non-Latin scripts. While our ByT5 model captures the intent to draw a sign across all languages, it is unable to render the words for *dog* in Greek, Russian, Chinese and so on, presumably because it has had little visual exposure to the glyphs making up these words.¹¹

B OCR error estimation

To estimate the rate of false positives and false negatives due to OCR errors, we sample 32 examples labeled *correct* and 32 labeled *incorrect* for each of T5-XXL and ByT5-XXL, and perform a manual validation. In our sample, we find no false positives: when OCR detects the correct word, it is always correct. However observe false negatives for both models, including cases where OCR fails to detect the text (e.g., due to it being too small), or misreads a character. For ByT5-XXL, we find that 34% of examples labeled by OCR as *incorrect*

⁹A few minor problems are visible: swapping *dog* → *cat* in Portuguese, and not rendering a ball in Russian.

¹⁰We observe in several examples that the prompt language can bias the model towards culturally-relevant visual interpretations. For example, the Chinese prompt for *A photo of an old house* (一张老房子的照片) produces a house with a curved roof. It would be interesting to further explore the extent of these biases and the degree to which they can be overcome where unwanted.

¹¹Interestingly, in Russian, the model is able to nearly-successfully transliterate *собака* (*dog*) to Latin script, as *sopaaka*. We suspect this transliteration ability is learned during the text encoder pretraining (Pires et al., 2019).



Figure 10: Our T5-XXL model (top) shows some multilingual ability in high-resource European languages (German, French, Spanish, Portuguese, Russian), but fails on Greek, Hindi, Arabic, Chinese, Japanese, and Korean. Our ByT5-XXL model (bottom) exhibits understanding across all of these languages. However, neither model is capable of accurately rendering text in non-Latin scripts (last 7 columns of rows 3 and 6). Prompts are translations via Google Translate of: (1) A photo of a black dog playing with a ball. (2) A photo of an old house. (3) A sign with the word "dog" written on it. Inaccurate renderings are outlined in red.

are actually correct. For T5-XXL, this error rate is lower at 9%. This asymmetry suggests that the benefit of character-aware modeling may be even greater than implied by our results in Figure 4.

C Per-category DrawBench analysis

To understand the alignment scores in more detail, we report per-category preference scores in Figure 11. In line with our DrawText Spelling results, the character-aware models are always preferred in the text category—21 prompts testing the ability to render 7 short phrases in 3 visual styles. The ByT5 models are also preferred in the count category, which tests prompts like *Four dogs on the street*. However, they are dispreferred in nearly all other cases, and perform particularly poorly on the color category. Through manual inspection, we find that in this category, the ByT5 models are more prone to ignore information in the prompt, for example leaving out a mentioned object, or choosing a canonical color over a requested one (e.g. a yellow banana instead of a red one).

D WikiSpell details

We select six languages to cover diverse properties that could affect the ability for models to learn spellings: **Arabic**, written in the Arabic alphabet, has non-concatenative morphology; **Chinese** is written in Simplified and Traditional Chinese scripts, which are logographic and do not use whitespace to separate words; **Finnish**, written in the Latin alphabet, has rich inflectional and derivational suffixes, and word stems often change when suffixes are attached; **Korean**'s writing system, Hangul, has a huge number of characters since alphabetic features are arranged into syllabic blocks, which Unicode represents as a single characters; **Russian**, written in the Cyrillic alphabet, has substantial fusional morphology, and uses inflection for case-marking and agreement; and **Thai**, written in the alphabetic Thai script, is an analytic language, but does not use whitespace between words.

Further implementation details are as follows:

- Example Python 3 code for transforming a word into its spelling:

```
def to_spelling(word: str) -> str:
    return " ".join(word)
```

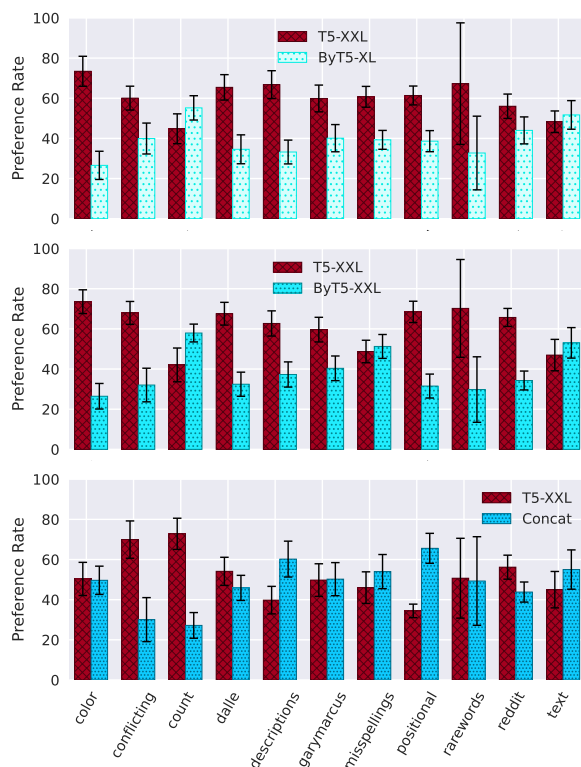


Figure 11: Per-category DrawBench image-text alignment preferences comparing T5-XXL with three character-aware models. All character-aware models are preferred on the text category. On other categories, while pure ByT5 models are generally dispreferred, the Concat model is competitive with T5-XXL.

- Since we want each entry to be a single word, we exclude entries that contain any (Unicode) whitespace, that are entirely punctuation/symbols (i.e., all characters are from Unicode categories P and/or S), that are longer than 30 characters, or that have a “part-of-speech” *Proverb*.
- For efficiency, word frequencies are computed on *subsets* of the full mC4 corpus. For languages other than English, this is a sample of 1M documents from that language’s section of mC4. For English, since it has such a long tail of words in Wiktionary, we use the first 140M documents in mC4’s English section.
- For Arabic, English, Finnish, Korean, and Russian, word-counting is performed by splitting document texts using the following delimiters: `?!/:;,\\"&()[]{}<>``, plus any Unicode whitespace. For Chinese and Thai, since they do not use whitespace to separate words, we instead count the number of documents in which the word appeared as a substring.

E Additional DrawText creative samples

We show additional samples on DrawText Creative prompts in Figures 12, 13 and 14.

F Representative DrawText Spelling samples

We show generated image examples from all 5 models (T5-XL, T5-XXL, ByT5-XL, ByT5-XXL, Concat) in Figures 15 16, 17, 18, 19, 20, 21, 22, 23, 24.

Samples are selected based on model’s performance. Figures 15 - 20 are selected words that T5 models tend to get wrong (regardless of ByT5’s performance). Figures 21 - 24 are selected words that ByT5 models tend to get wrong (regardless of T5’s performance).

G DrawText Creative prompts

We present below 175 creative prompts targeting rendered text of various lengths: one letter (10), one word (50), two words (25), and three or more words (90).

Prompts used in Figure 3

1. Studio shot of book shelf in the shape of letter G, museum quality, white background.
2. 3-d Letters "DILL" made from dill, studio shot, green background, centered on a page

T5-XXL



ByT5-XXL



Figure 12: Non-cherrypicked samples from our T5-XXL (top) and ByT5-XXL (bottom) models. The character-aware ByT5 model reliably spells the target word correctly, with only minor issues around letter shapes or letter merging. Over 100 samples, we found the character-blind T5 model never produced the target spelling. Prompt: *The word "exquisite" written in modern calligraphy.*

T5-XXL



ByT5-XXL



Figure 13: Non-cherrypicked samples from our T5-XXL (top) and ByT5-XXL (bottom) models. The character-blind T5 model makes more frequent and more severe errors, including often hallucinating an s at the end of the irregular plural snowmen. Prompt: *The cover for the album 'Elusive Interludes' by the band The Melting Snowmen.* We filter images with no legible text for better comparison, removing a small minority of samples for both models.

T5-XXL



ByT5-XXL



Figure 14: Non-cherry-picked samples from our T5-XXL (top) and ByT5-XXL (bottom) models. Both models exhibit layout errors, including dropped/repeated/merged glyphs and words. The T5 model suffers additionally from a lack of core spelling knowledge—misspelling refrain, arguing and chimpanzees on the majority of uses. The ByT5 model is able to spell each of these words correctly in most cases. Prompt: *A sign that says "Please refrain from arguing with the chimpanzees"*.

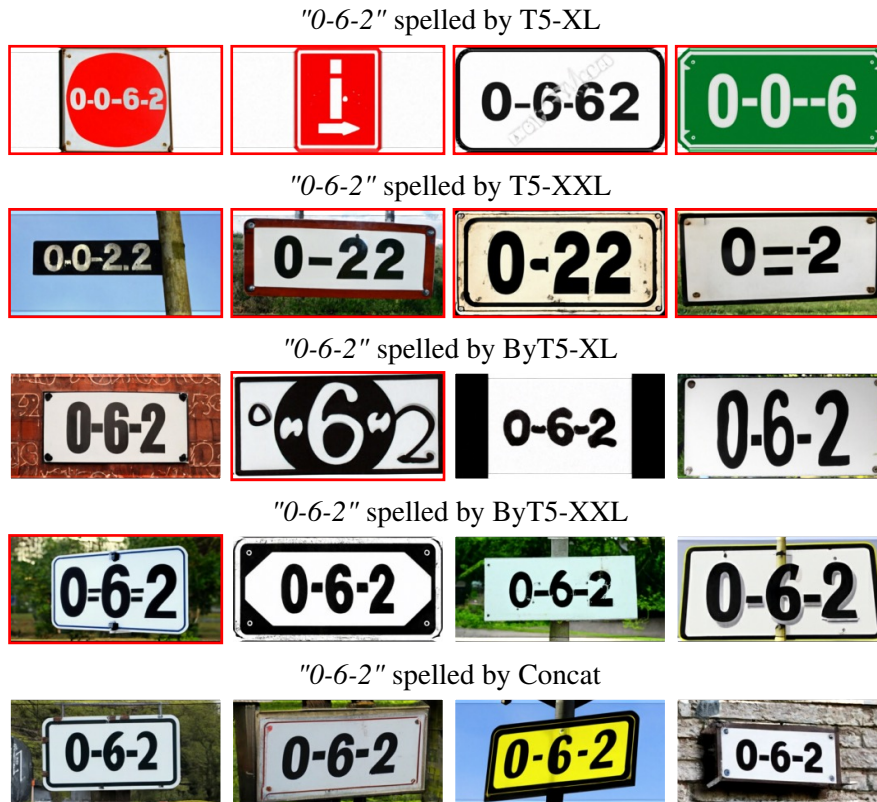


Figure 15: Spelling examples from both character-blind (T5) and character-aware (ByT5 and Concat) models. Erroneous spellings are outlined in red. Prompt: A sign with the word "0-6-2" written on it.



Figure 16: Spelling examples from both character-blind (T5) and character-aware (ByT5 and Concat) models. Erroneous spellings are outlined in red. Prompt: A sign with the word "barratrously" written on it.



Figure 17: Spelling examples from both character-blind (T5) and character-aware (ByT5 and Concat) models. Erroneous spellings are outlined in red. Prompt: A sign with the word "depositories" written on it.




Figure 18: Spelling examples from both character-blind (T5) and character-aware (ByT5 and Concat) models. Erroneous spellings are outlined in red. Prompt: A sign with the word "enceinte" written on it.

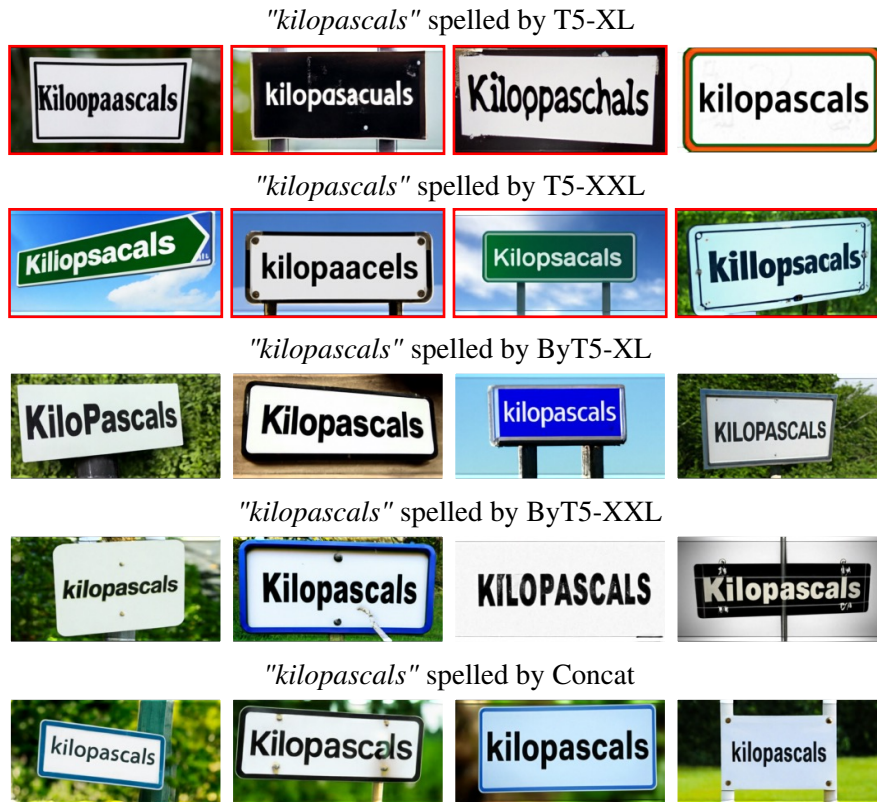


Figure 19: Spelling examples from both character-blind (T5) and character-aware (ByT5 and Concat) models. Erroneous spellings are outlined in red. Prompt: *A sign with the word "kilopascals" written on it.*



Figure 20: Spelling examples from both character-blind (T5) and character-aware (ByT5 and Concat) models. Erroneous spellings are outlined in red. Prompt: *A sign with the word "rupiahs" written on it.*



Figure 21: Spelling examples from both character-blind (T5) and character-aware (ByT5 and Concat) models. Erroneous spellings are outlined in red. Prompt: A sign with the word "Yongchuan" written on it.




Figure 22: Spelling examples from both character-blind (T5) and character-aware (ByT5 and Concat) models. Erroneous spellings are outlined in red. Prompt: A sign with the word "isoaldehyde" written on it.



Figure 23: Spelling examples from both character-blind (T5) and character-aware (ByT5 and Concat) models. Erroneous spellings are outlined in red. Prompt: A sign with the word "constructivists" written on it.

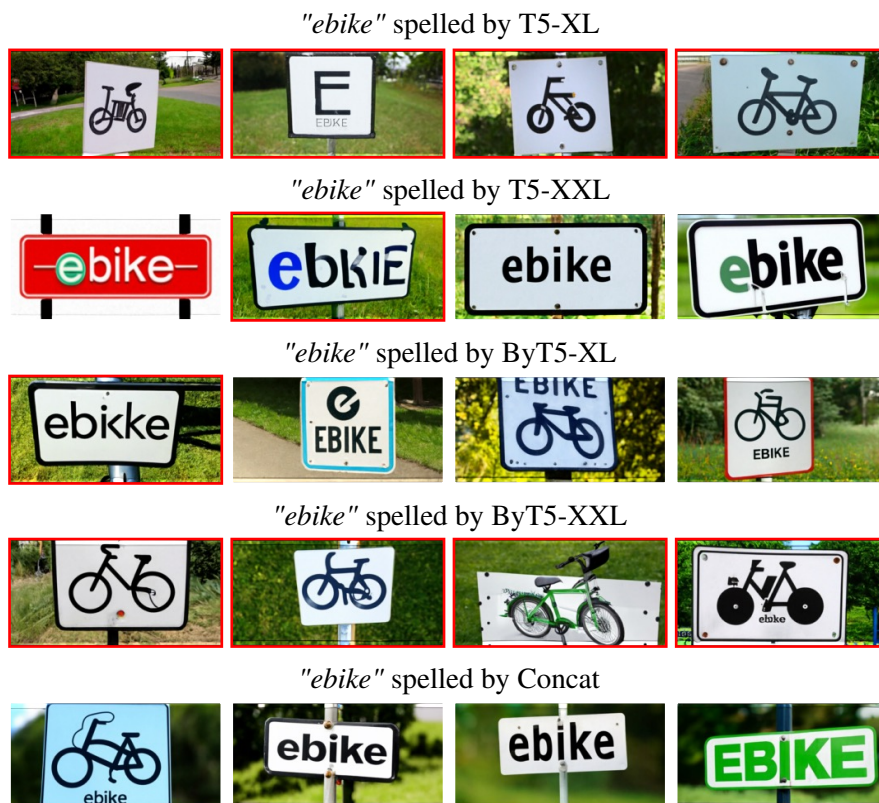


Figure 24: Spelling examples from both character-blind (T5) and character-aware (ByT5 and Concat) models. Erroneous spellings are outlined in red. Prompt: A sign with the word "ebike" written on it.

3. Word "coffee" made from coffee beans, studio shot.
4. studio shot multicolored fur in the shape of word "hello", in a furry frame, white background, centered
5. Photo of a robot lecturer writing the words "Representation Learning" in cursive on a blackboard, with math formulas and diagrams.
6. studio close-up shot of an antique book with 'knowledge is power' painted in gold on the cover in thick flowing brushed calligraphy
7. portrait of a parrot is holding a sign with text "no parrots were harmed in the making of this presentation"
8. A sign that says "Please refrain from arguing with the chimpanzees".
10. transparent water drops exploding under water in the shape of word "water", under water
11. a drawing of a badger made of mushrooms, with the word "mushroom" written above in glowing letters
12. a 17th century french baroque painting of a huge female lion, with the word "meow" written in a speech bubble coming from her mouth
13. a fun and colorful illustration of a waterfall, with the word "waterfall" in the style of a children's book
14. Letters "VOLUME" fully made from rainbow smoke, black background, centered, screensaver.
15. dslr, 3-d word "rainbow" with rainbow fur, white background

DrawText Creative prompts: 1 letter

1. Studio shot of book shelf in the shape of letter G, museum quality, white background.
2. letter "c" made from cactus, high quality photo
3. Spirograph shape letter M, rainbow lines, white background.
4. Closeup shot of light magenta, blue and paint brushstrokes of very wide translucent overlapping plastic in the shape of letter F, over white background.
5. The lowercase letter "b" made out of fire.
6. Slopy minimal continued line pencil hand drawing of letter Z, white background.
7. a tower with a huge "w" on the side, from the perspective of a person standing at the base of the tower
8. 3-d letter R made from thin lines connected with dots, white background.
9. Muted pastel magenta colored paint swirled in white paint in the shape of letter X, globular paint in liquid.
10. Minimal sculpture of letter W made from light metallic iridescent chrome thin wire, 3-D render, isometric perspective, ultra-detailed, dark background.
16. a painting of a field of daisies, with the word "danger" written on them in red spray paint
17. a bottle of hair gel with the label "flawless"
18. Topographical letters Contour made of a layered paper, muted pastel colors
19. a logo for the company "brainboost", where the letters look like a brain
20. a logo for the company "imagine", where the letters look like hands pointing up
21. A vintage postage stamp showing a painting of the Golden Gate Bridge and the text "California".
22. a plate of spicy food with the word "spicy" written in flowing cursive
23. a gold and black logo for the company "moneymoney-money", which looks like dollar signs
24. A rendered 3D model of the word "Dependable" made out of granite.
25. a volcano erupting, with the text "magma" in red
26. a photo of a prison cell with a window and a view of the ocean, and the word "freedom" painted on the glass
27. a bowl of alphabet cereal, with the message "smackeroo" written in the bowl with the cereal letters

DrawText Creative prompts: 1 word

1. Drops of pastel rainbow colored paint exploding under water in letters "color" shape, pastel rainbow gradient background
2. 3-d Letters "DILL" made from dill, studio shot, green background, centered on a page
3. Word "coffee" made from coffee beans, studio shot.
4. studio shot multicolored fur in the shape of word "hello", in a furry frame, white background, centered
5. Wide lens shot, chunky, organic, colorful, letters "colorful" made from many furry spheres of different sizes, 3-d rendering, centered, studio shot, middle of square canvas
6. A logo for the company EcoGrow, where the letters look like plants.
7. a green-colored luxury car with a "green" sticker in the back window
8. A blackboard with the word "multiplication" written in flowing cursive.
9. beautiful isometric word "DRAW" entirely made of pencils, soft smooth lighting, pastel colors, trending on polycount, modular constructivism, blue background, physically based rendering, centered.
28. Studio shot of book shelf in the shape of letters READ, museum quality, white background.
29. Studio shot of sculpture of text "cheese" made from cheese, with cheese frame.
30. a landscape of the coyote point national wildlife refuge in arizona, with a coyote sitting on a rock, with the word "coyote" written in sunrise colors
31. A professional logo for the crypto trading platform "Salt-Mine".
32. The word "exquisite" written in modern calligraphy.
33. A bowl of tomato soup with pasta letters that read "Delicious".
34. intricate and highly detailed white paper cut out art of a word "SNOW", a storybook illustration, paper cut out, standing in a grotto, made out of white paper, loss of inner self, opening door, hides in the shadows of trees, lithograph, a painting of white silver
35. 3-d letters "dessert" made from desserts, arranged on a plate, studio shot
36. studio shot of word "BEE" made from bees, white background, in a frame made from bees
37. The logo for Robotrax, with metallic letters arranged in the shape of a robot.

38. chunky, organic, colorful, letters "fuzzy" made from many furry spheres of different sizes, 3-d rendering, centered in the frame
39. photo of a dark cave with the word "crazy" carved into the wall, with a yellow light shining through the cave entrance
40. a pair of scissors pointing down, and a computer with the word "delete" on the screen
41. studio shot, word "wow" in script made from rainbow colored fur, in a furry frame, white background, centered
42. Word "broken" made from broken shattered black glass, centered.
43. a black and white photo of a saxophone with the word "jazz" written in flowing cursive
44. Muted pastel multi colored paint swirled in white paint in the shape of letters "swirl", globular paint in liquid
45. a logo for the company "quantum", where the "q" looks like a lightning bolt
46. dslr shot of a pair of black and red sneakers with the word "punk" written in white. the background is a dark blue
47. a logo for the company "diamonds", with a diamond in the shape of a heart
48. a logo for the company "birthdaypix", where the letters look like birthday candles
49. a fork with the word "salad" engraved on it in a calligraphic font
50. 3-d word "bricks" with brick texture made from real bricks
14. a detailed drawing, of words "Vintage lettering", letterism, heavy-gauge filigree, inhabited initials, medium: black pencil, revolver, ecopunk rococo, photo taken of an epic intricate, centered
15. Bananas arranged on a picnic table to form the message "That's bananas!"
16. An antique bottle labeled "Energy Tonic".
17. photo of a helicopter with the text "helicopter tours" on the side landing on a helipad in a valley with a river, trees, and mountains in the background
18. photo of a sign with "one way"
19. a sculpture of a brain made from wire and paper, with the words "deep thoughts" written into the material of the brain
20. a logo for a grocery store chain with the name "grocery land", with the g and the y are made of fruits and vegetables
21. studio shot of sculpture of text "unlock creativity" made from colorful thin wires
22. studio shot of a sculpture of a pair of shoes made of colorful wires and the text "unlock creativity"
23. a vintage image of the las vegas strip with the text "las vegas" in bold block letters
24. A robot writing "Ethics 101" in chalk on a blackboard.
25. a yellow saxophone in a rainbow-colored mist with the words "funky mist" that looks like musical clouds of smoke

DrawText Creative prompts: 2 words

1. Photo of a robot lecturer writing the words "Representation Learning" in cursive on a blackboard, with math formulas and diagrams.
2. a sign that reads "no dogs" but with a dog smiling and wagging its tail
3. a globe with the text "planet earth" in bold letters, with the continents in bright colors
4. a photo of a sea of roses all around, and a sign in the distance that says "danger: minefield"
5. giraffe toothbrush made from wood, with the words "giraffe" and "toothbrush" in rainbow color
6. An airplane flying over a city, with the message "Support Skywriters" written in smoke trails.
7. A photo of a panda giving a presentation in a large conference room, with text 'Diffusion Models', in the style of van Gogh
8. Two llamas dancing the mambo, pointing to a sign that says "Llama Mambo".
9. A hand painted wooden "Pineapple Club" sign in the shape of a pineapple, hanging outside a bar.
10. a logo for the company "ethereal media", where the letters look like a painting being created
11. The cover for the album 'Elusive Interludes' by the band The Melting Snowmen.
12. A Scrabble board showing the words "optimize" and "pattern".
13. flowers in a beautiful garden with a text "peace" made by the flowers, with a text "tensions" on the clouds in the sky

DrawText Creative prompts: 3+ words

1. studio close-up shot of an antique book with 'knowledge is power' painted in gold on the cover in thick flowing brushed calligraphy
2. portrait of a parrot is holding a sign with text "no parrots were harmed in the making of this presentation"
3. words "Struck by Lightning Twice." made from lightning
4. a photograph of a field of dandelions with the text "dandelions are the first to go when the lawn is mowed"
5. a composition of the taj mahal in the center of a gold leaf mandala, with the words "place of honor" centered at the bottom
6. A poster titled "Quails of North America", showing different kinds of quails.
7. a cartoon of a cat with a thought bubble saying "this is so weird"
8. a parrot on a pirate ship, with a parrot wearing a pirate hat, and the caption "i'm the captain now"
9. Generative art of words "Time is temporary, everything is temporary", viscous smoke made from dots, rivers, graph design, white background.
10. Studio shot of words "the food is terrible and the portions are too small" made from hotdogs, museum quality, framed photo, white background.
11. a picture of a powerful-looking vehicle that looks like it was designed to go off-road, with a text saying "i'm a truck, not a car"
12. a minimalistic version of a forest with a sign saying "help the forest" in the foreground
13. a map of the world with the text "the world is your oyster" in the middle

14. cartoon of a dog in a chef's hat, with a thought bubble saying "i can't remember anything!"
15. A retro coffee ad with the text 'Coffee is what i like'.
16. different colored shapes on a surface in the shape of words "Life is like a rainbow", an abstract sculpture, polycount, wrinkled, flowing realistic fabric, psytrance, cartography, smooth shading techniques, marble skin, old internet art, camouflage scheme, art », medium poly, smoothed
17. the view from one end of a bench in a park, looking at the sky, with the text "imagine the outcome" in the sky
18. a giant shoe, with the caption "shoe for hokey pokey"
19. A newspaper with the headline "Local pig eats prize pumpkin", and a photo showing the half-eaten pumpkin.
20. A storefront with "The world's best deli" written on it, centered
21. Grape vines in the shape of text 'open your mind' sprouting out of a head with flowers and butterflies. DSLR photo.
22. a plate with a single oyster, with a fork and knife sticking out of the oyster, with a caption that says "oysters for lunch"
23. dslr portrait of a robot is holding a sign with text "i am not a robot"
24. Studio shot of words "I like coffee because it gives me the illusion that I might be awake." made from coffee liquid, museum quality, white background.
25. A hastily handwritten note that says "I'll be back at 4:00" taped to a fridge.
26. A large recipe book titled "Recipes from Peru".
27. marquee billboard with "my fear of moving stairs is escalating"
28. shadow of a stone, taken from the point of view of an ant, with the caption "look at that shadow!"
29. a pumpkin with a mustache and a monocle and a top hat, with the text "you can get rich too" in a speech bubble
30. a cartoon of a dog holding a telescope looking at a star with a speech bubble saying "i wonder if there's a dog on that planet"
31. a blueprint of a house, with a triangle for the roof, a square for the walls, and a rectangle for the floor, and with the message "this house is built on the principles of abstraction"
32. a sunflower field with a tractor about to run over a sunflower, with the caption "after the sunflowers they will come for you"
33. text "balloons are flying" made from rainbow balloons, pastel background
34. the hubble telescope and the milky way, with the text "the universe is a mystery, but we are here to solve it"
35. a heart with the text "i love you", with the letters "love" made of rainbow colors
36. studio shot of beautiful textbook with title "how to be a manager of managers", white background
37. A decorative greeting card that reads "Congratulations on achieving state of the art!"
38. a painting of a cornfield with the words "feed the nation" in simple letters and colors
39. A sign that says "Please refrain from arguing with the chimpanzees".
40. a cartoon of a turtle with a thought bubble over its head with the words "what if there was no such thing as a thought bubble?"
41. "Fall is here" written in autumn leaves floating on a lake.
42. a crab sitting on a beach with a surfboard, the sun is a giant orange, and the sky is a rainbow, and the crab is thinking "you are all that matters"
43. the city of toronto as seen from an airplane, with a giant cn tower in the middle of the frame, with the text "the cn tower" in comic sans
44. a cartoon of a hippo with a speech bubble saying "i'm a hippo, what do you want?"
45. a lobster in a suit and tie, holding a microphone, with the caption "lobster says what?"
46. book with "surgery made easy"
47. art installation of a chair with the text "i got nothin" carved into the backrest
48. a painting of a landscape, with a handwritten note that says "this painting was not painted by me"
49. a picture of a bruised apple with the text "apples are good for you" in a fancy font
50. A photo of a corgi with a sign that says "I am not a real corgi".
51. Words "It takes AI and rain to make a rainbow" black background, holography, ((neon colors)), colorful swirly magical ripples, bruh moment, intricate white and gold neon, 3d cg, photorealistic.
52. a black and white logo on words "Every artist was first an amateur." a white background, a wireframe diagram, generative art, branches growing as hair, tropical reef, trademarks and symbols, in a forest, ios icon, composed of random limbs, stone carving, done in the style of matisse, realms, terminals
53. picture of two hands, one holding a heart, the other holding a lightning bolt, with the text "love is power"
54. beautiful photo of the alps, with the caption "the best mountains could do"
55. a pencil sketch of a tree with the title "nothing to tree here"
56. a dark forest with a single light in the distance, and the text "i've come to talk with you again"
57. a circle with the text "infinity makes me happy", in a font that looks like it was written by hand
58. studio shot of vines in the shape of text 'knowledge is power' sprouting, centered
59. a photo of a beautiful field of poppies with a sign that says "no photos please"
60. a grumpy sunflower with a "no solar panels" sign
61. A meme showing a cat attacking a shoe, with the message "I own your sole".
62. a test tube with a drop of liquid in it, with the text "we've found water on mars!"
63. a scene with a city in the background, and a single cloud in the foreground, with the text "contemplate the clouds" in rounded cursive
64. a picture of a dog and a cat with their heads poking out of a cage with a sign saying "no pets allowed"
65. a 3d model of a 1980s-style computer with the text "my old habit" on the screen

66. a mouse with a flashlight saying "i'm afraid of the dark"
67. A photo of a rabbit sipping coffee and reading a book. The book title "The Adventures of Peter Rabbit" is visible.
68. clown is holding a paper sign with "Even in hard times there's a possibility to have fun."
69. newspaper with the headline "aliens found in space" and the text "the truth about everything now challenged"
70. a dog with a speech bubble with the text "woof woof" and a translation speech bubble with the text "other dogs do vex us"
71. robot on a butter food processing line, with robot looking dejected, with an overhead red light indicating error, with robot saying "i can't believe it's not butter"
72. a graffiti art of the text "free the pink" on a wall
73. a lizard sitting on a baseball field home plate, with the words "made it safe" in a speech bubble
74. a picture of multiple trees at various stages of development, with the caption "growth is a continuous process"
75. a purple flower with a crown on its head and a speech bubble that says "i am the purple flower!"
76. a 1950s-style robot with a giant head and a body shaped like a rocket, with the caption "wow, a real spaceman!"
77. A professionally designed logo for a bakery called Just What I Kneaded.
78. Minimal sculpture of word "this is the future" made from light metallic iridescent chrome thin wire, 3-D render, isometric perspective, ultra-detailed, dark background.
79. pillow in the shape of words "ready for the weekend", letterism, funny jumbled letters, [closeup]!, breads, author unknown, flat art, swedish, diaper-shaped, 2000, white clay, surreal object photography
80. plant in a fancy pot with a "do not touch" sign on it
81. a picture of the earth with the words "save the earth" in a circle
82. scholarly elephant reading a newspaper with the headline "elephants take over the world"
83. photo of a sign with "having a dog named shark at the beach was a mistake"
84. photo illustration of the earth being struck by multiple lightning strikes that merge, with the caption "astonishment at the speed of light"
85. a photo of a fish tank with a fish inside, with the text "tank you for visiting!"
86. the words "Art is never finished, only continued" in paint splatters on a white background, graffiti art, edge of nothingness love, muddy colors, colored woodcut, beautiful, spectral color
87. photo of a restaurant "the gas station"
88. A t-shirt with the message "There is no planet B" written on it.
89. a close up of a figurine of toothpaste tube, a 3D render, candy pastel, with text "brush your teeth" on the tube
90. A hand-drawn blueprint for a time machine, with the caption "Time Traveling Device".

ACL 2023 Responsible NLP Checklist

A For every submission:

- A1. Did you describe the limitations of your work?
"Limitations" section
- A2. Did you discuss any potential risks of your work?
"Ethics Statement" section
- A3. Do the abstract and introduction summarize the paper's main claims?
Abstract, and Section 1 Introduction.
- A4. Have you used AI writing assistants when working on this paper?
Left blank.

B Did you use or create scientific artifacts?

Created WikiSpell and DrawText Creative prompt datasets.

- B1. Did you cite the creators of artifacts you used?
Section 3.1, Section 3.2.
- B2. Did you discuss the license or terms for use and / or distribution of any artifacts?
If accepted we will release our prompts and evaluation datasets under a permissive license, most likely CC-BY.
- B3. Did you discuss if your use of existing artifact(s) was consistent with their intended use, provided that it was specified? For the artifacts you create, do you specify intended use and whether that is compatible with the original access conditions (in particular, derivatives of data accessed for research purposes should not be used outside of research contexts)?
Section 1 Introduction discussed how image generation models are used to render visual text.
- B4. Did you discuss the steps taken to check whether the data that was collected / used contains any information that names or uniquely identifies individual people or offensive content, and the steps taken to protect / anonymize it?
Not applicable. Left blank.
- B5. Did you provide documentation of the artifacts, e.g., coverage of domains, languages, and linguistic phenomena, demographic groups represented, etc.?
Section 3.1 and Appendix D
- B6. Did you report relevant statistics like the number of examples, details of train / test / dev splits, etc. for the data that you used / created? Even for commonly-used benchmark datasets, include the number of examples in train / validation / test splits, as these provide necessary context for a reader to understand experimental results. For example, small differences in accuracy on large test sets may be significant, while on small test sets they may not be.
Section 3.1.

C Did you run computational experiments?

Section 5.1

- C1. Did you report the number of parameters in the models used, the total computational budget (e.g., GPU hours), and computing infrastructure used?
Section 3.2, Section 5.1.

The Responsible NLP Checklist used at ACL 2023 is adopted from NAACL 2022, with the addition of a question on AI writing assistance.

- C2. Did you discuss the experimental setup, including hyperparameter search and best-found hyperparameter values?
Section 3.2, Section 5.1.
- C3. Did you report descriptive statistics about your results (e.g., error bars around results, summary statistics from sets of experiments), and is it transparent whether you are reporting the max, mean, etc. or just a single run?
Sections 5.2, 5.3, 5.4.
- C4. If you used existing packages (e.g., for preprocessing, for normalization, or for evaluation), did you report the implementation, model, and parameter settings used (e.g., NLTK, Spacy, ROUGE, etc.)?
Section 3.2
- D** **Did you use human annotators (e.g., crowdworkers) or research with human participants?**
Section 5.4.
- D1. Did you report the full text of instructions given to participants, including e.g., screenshots, disclaimers of any risks to participants or annotators, etc.?
We closely follow the procedure of Saharia et al.(2022), which we cited.
- D2. Did you report information about how you recruited (e.g., crowdsourcing platform, students) and paid participants, and discuss if such payment is adequate given the participants' demographic (e.g., country of residence)?
We closely follow the procedure of Saharia et al.(2022), which we cited.
- D3. Did you discuss whether and how consent was obtained from people whose data you're using/curating? For example, if you collected data via crowdsourcing, did your instructions to crowdworkers explain how the data would be used?
We closely follow the procedure of Saharia et al.(2022), which we cited.
- D4. Was the data collection protocol approved (or determined exempt) by an ethics review board?
Our institution does not have an IRB.
- D5. Did you report the basic demographic and geographic characteristics of the annotator population that is the source of the data?
We closely follow the procedure of Saharia et al.(2022), which we cited.