

Trick or : Thematic Reinforcement for Artistic Typography

Purva Tendulkar¹ Kalpesh Krishna² Ramprasaath R. Selvaraju¹ Devi Parikh¹

¹Georgia Institute of Technology, ²University of Massachusetts, Amherst
¹{purva, ramprs, parikh}@gatech.edu, ²kalpesh@cs.umass.edu

Abstract

An approach to make text visually appealing and memorable is *semantic reinforcement* – the use of visual cues alluding to the context or theme in which the word is being used to reinforce the message (e.g., Google Doodles). We present a computational approach for semantic reinforcement called TReAT – Thematic Reinforcement for Artistic Typography. Given an input word (e.g. *exam*) and a theme (e.g. *education*), the individual letters of the input word are replaced by cliparts relevant to the theme which visually resemble the letters – adding creative context to the otherwise ordinary and uninspiring input word. We use an unsupervised approach to learn a latent space to represent letters and cliparts and compute similarities between the two. Human studies show that participants can reliably recognize the word as well as the theme in our outputs (TReATs) and find them more creative compared to meaningful baselines.

Introduction

We address the task of theme-based word typography: given a word (e.g., *exam*) and a theme (e.g., *education*), the task is to *automagically* produce a “doodle”¹ for the word in that theme as seen in Figure 1. Concretely, the task is to replace each letter in the input word with a clipart from the input theme to produce a doodle, such that the word and theme can be easily identified from the doodle. Solving this task would be of value to a variety of creative applications such as stylizing text in advertising, designing logos – essentially any application where a message needs to be conveyed to an audience in an effective and concise manner.

Using graphic elements to emphasize the meaning of a word in reference to a related theme is referred to in graphic design as *semantic reinforcement*. This can be achieved in a number of ways, e.g., using different fonts and colors (Figure 2a), changing the position of letters relative to one another (Figure 2b), arranging letters in a specific direction or shape (Figure 2c), excluding some letters (Figure 2d), adding icons near or around the letters (Figure 2e), or replacing letters with icons (Figure 2f). In our work, we focus on this last type, i.e., semantic reinforcement via replacement.

¹In this work, we use “doodle” to refer to Google doodles-like typography as in Figure 2f.

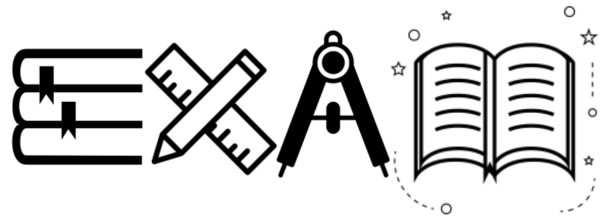


Figure 1: A sample doodle (that we call TReAT) generated by our system for the input word *exam* and theme *education*².

This is a challenging task even for humans. It not only requires domain-specific knowledge for identifying a set of relevant cliparts to choose from, but also requires creative abilities to be able to visualize a letter in a clipart, and choose the best clipart for representing it.

The latter alone is challenging to automate – both from a training and evaluation perspective. Training a model to automatically match letters to graphics is challenging because there is a lack of large-scale text-graphic paired datasets in each domain that might be of interest (e.g., clipart, logogram). Evaluation and thus iterative development of such models is also challenging because of subjectivity and inter-human disagreement on which letter resembles which graphic.

We present a computational approach – Thematic Reinforcement of Artistic Typography (TReAT) – to generate doodles (TReATs) for semantic reinforcement of text via replacement. We represent letters in different fonts and cliparts from the Noun Project. in a common latent space. These latent representations are learned such that they have two characteristics: (1) The letters can be correctly recognized (e.g., *a* vs. *b*) in the latent space and (2) The letters and cliparts can be reconstructed accurately from the latent space. A reconstruction loss ensures that letters and clipart that are close in the latent space also look similar in the image space. A classification loss ensures that the latent space is informed by discriminative features that make one letter different from the other. This allows us to match cliparts to letters in a way that preserves distinctive visual features of the letters, making it easier for humans to identify the letter being depicted by a

²Unless stated otherwise, all cliparts in the paper have been taken from The Noun Project - <https://thenounproject.com/>. The Noun Project contains cliparts created by different graphic designers on a variety of themes.




Figure 2: Different methods for semantic reinforcement. a) font and color variations b) positioning of letters relative to each other c) arrangement of letters in a specific shape or direction d) exclusion of some letters e) addition of icons near letters f) replacement of letters. In this work we focus on f), semantic reinforcement via replacement.³

clipart. At test time, given a word and a theme as input, we first retrieve cliparts from the Noun Project that match that theme. For each letter in the word, we find the theme-relevant clipart which minimizes the distance from it across a variety of fonts. If the distance is low enough, we replace the letter with the clipart.

We run human studies to show that subjects can reliably recognize the word as well as the theme from our TReATs, and find them to be creative. We consider a TReAT to be creative if it is surprising and/or intriguing and/or fun.

Related work

Early human communication was through symbols and hieroglyphs (Frutiger 1989), (Schmandt-Besserat 2015). This involved the use of characters to represent an entire word, phrase or concept. Then language evolved and we started using the alphabet for creation of new words to represent concepts. However many languages (e.g. Chinese and Japanese) still make use of pictograms or logograms to depict specific words. Today, symbols and logos are used for creative applications to increase the communication bandwidth – to convey abstract concepts, express rich emotions, or reinforce messages (Shiojiri and Nakatani 2013), (Clawson et al. 2012), (Takasaki and Mori 2007). Our work produces a visual depiction of text by reasoning about similarity between the visual appearance of a letter and clipart imagery. We describe prior work in each of these domains: creativity through imagery and creativity through visual appearance of text.

Creativity through imagery

There has been work on evoking emotional responses through the modification of images. Work on visual blending of emojis combines different concepts to create novel emojis (Cunha et al. 2018). Visual blending has also been explored for combining two animals to create images depicting fictional hybrid animals (Martins et al. 2015). Our approach tries to induce creativity by entirely replacing a letter with a clipart.

Work on Vismantic (Xiao and Linkola 2015) represents abstract concepts visually by combining images using juxtaposition, fusion, and replacement. Our work also represents a theme via replacement (replacing letters with cliparts); however our replacement is for the purposes of lexical resolution, not visual. Recently, there has been an exploration of neural style transfer for logo generation (Atarsaikhan, Iwana, and Uchida 2018). This work however only transfers color and texture from the style to the content. In our work, the

transfer occurs via direct replacement. Logo generation has also been explored through the use of Generative Adversarial Networks (Sage et al. 2018).

Recently Google’s QuickDraw! and AutoDraw based on sketch-rnn (Ha and Eck 2018) have gained a lot popularity. Their work trains a recurrent neural network (RNN) to construct stroke-based drawings of common objects, and is also able to recognize objects from human strokes. One could envision creating a doodle by writing out one letter at a time, that AutoDraw would match to the closest object in its library. However, these matches would not be theme based. Iconary⁴ is a very recent pictiary-like game that users can play with an AI. Relevant to this work, user drawings in Iconary are mapped to icons from The Noun Project to create a scene.

The use of conditional adversarial networks for image-to-image translation is gaining popularity. However, using a pix2pix-like architecture (Isola et al. 2017) for our task would involve the use of labeled pairwise (clipart, letter) data, which as discussed earlier, is hard to obtain. CycleGAN (Zhu et al. 2017) does not require paired label data, but is not a good fit for our task because we are interested in matching letters to cliparts from a specific theme. The pool of clipart is thus limited, and would not be sufficient to learn the target domain. Finally, generative modeling is typically lossy; we prefer direct replacement of cliparts for greater readability.

Creativity through visual appearance of text

Advances in conditional Generative Adversarial Networks (cGANs) (Mirza and Osindero 2014) have motivated style transfer for fonts (Azadi et al. 2018) through few-shot learning. Work on learning a manifold of fonts (Campbell and Kautz 2014) allows everyday users to create and edit fonts by smoothly interpolating between existing fonts. (Martins et al. 2018) explore an evolutionary system for the generation of type stencils constructed from line segments. The first work explores the creation of unseen letters of a known font, while the other two works explore the creation of entirely new fonts – neither add any theme-related semantics or additional graphic elements to the text.

Work on neural font style transfer between fonts (Atarsaikhan et al. 2017) explores the effects of using different weighted factors, character placements and orientations in style transfer. This work also has an experiment using icons as style images, however the style transfer is only within the context of visual features of icons such as the texture and thickness of strokes, as opposed to direct replacement.

³Examples a) to e) were taken from this answer on StackExchange. Example f) is a Google Doodle.

⁴<https://iconary.allenai.org/>

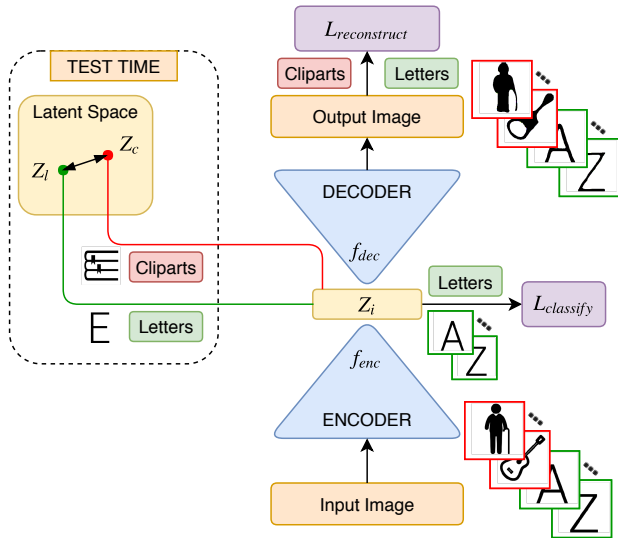


Figure 3: Block diagram describing our model during training and testing. The model is trained on a reconstruction and classification loss in a multitask fashion. During inference, latent space distances are calculated to match letters to cliparts. See text for more details.

MarkMaker⁵ generates logos based on company names – primarily displaying the name in various fonts and styles, sometimes along with a clipart. It uses a genetic algorithm to iteratively refine suggestions based on user feedback.

Approach

In this section, we first describe our procedure for collecting training data, and then our model and its training details. Finally, we describe our test-time procedure to generate a TReAT, that is, obtaining theme-based clipart matches for an input word and theme. A sketch of our model along with examples for training and testing are shown in Figure 3.

Training Data

For our task we need two types of data for training – letters in different fonts, and cliparts. Note that we do not need a *correspondence* between the letters and cliparts. In that sense, as stated earlier, our approach is an unsupervised one.

For clipart data, we use the Noun Project – a website that aggregates and categorizes symbols that are created and uploaded by graphic designers around the world. The Noun Project cliparts are all 200×200 in PNG format. The Noun Project has binary cliparts, and will result in TReATs of the style shown in Figure 1. Different choices of the source of cliparts can result in different styles, including colored TReATs similar to Figure 2f. We downloaded a random set of $\sim 50k$ cliparts from the Noun Project.

We obtain our letter data from a collection of 1400 distinct font files.⁶ On manual inspection, we found that this set contained a lot of visual redundancies (e.g. the same font being repeated in regular and bold weight types). We removed such repetitions. We also manually inspected the data to ensure

that the individual letters were recognizable in isolation, and discarded overly complicated and intricate font styles. This left us with a total of 777 distinct fonts. We generated 200×200 image files (PNG format) from each font file for the entire alphabet (uppercase and lowercase) giving us a total of 40.4k images of letters ($777 \text{ fonts} \times 26 \text{ letters in the English alphabet} \times 2 \text{ (upper and lower cases)}$).

Model

Our primary objective is to find visual similarities between cliparts and letters in an unsupervised manner. To this end, we train an autoencoder (Ballard 1987) with a reconstruction loss on both clipart and letter images (denoted by \mathbf{X}_{cl}). We denote a single input image by X_i . Each input image X_i is passed through an encoder neural network $f_{enc}(\cdot)$ and projected to a low dimensional intermediate representation Z_i . Finally, a decoder neural network $f_{dec}(\cdot)$ tries to reconstruct the input image as \hat{X}_i , using the objective $\mathcal{L}_{reconstruct}$,

$$Z_i = f_{enc}(X_i) \quad (1)$$

$$\hat{X}_i = f_{dec}(Z_i) \quad (2)$$

$$\mathcal{L}_{reconstruct} = \frac{1}{|\mathbf{X}_{cl}|} \sum_{i \in \mathbf{X}_{cl}} \text{SSD}(X_i, \hat{X}_i) \quad (3)$$

where $\text{SSD}(X_i, \hat{X}_i)$ is the sum over squared pixel differences between the original image and its reconstruction. We set the dimensionality of Z_i to be 128. In addition to the reconstruction objective, we utilize letter labels (52 labels for lowercase and uppercase letters) to classify the intermediate representations Z_i for the letter images. This objective helps the encoder discriminate between different letters (possibly with similar visual features) while clustering together the intermediate representations for the same letter across different fonts. This would allow the intermediate representation to capture visual features that are characteristic of each letter, and when cliparts are matched to letters using this representation, the matched cliparts will retain the visually discriminative features of letters.

Concretely, we project Z_i to a 52-dimensional space using a single linear layer with a softmax non-linearity and use the cross entropy loss function. Let W and b be the parameters of a linear transformation of Z_i . We obtain a probability distribution $P_i(\cdot)$ across all labels as,

$$P_i(\cdot) = \text{softmax}(WZ_i + b) \quad (4)$$

Let \mathbf{X}_l be the subset of images in \mathbf{X}_{cl} that are letters. We maximize the probability of the correct label Y_i corresponding to each letter image X_i .

$$\mathcal{L}_{classify} = -\frac{1}{|\mathbf{X}_l|} \sum_{i \in \mathbf{X}_l} \log P_i(Y_i) \quad (5)$$

Note that the same Z_i is used in both objective functions for letter images. These objectives are jointly trained using a multitask objective

$$\mathcal{L} = \alpha \mathcal{L}_{reconstruct} + (1 - \alpha) \mathcal{L}_{classify} \quad (6)$$

Our final loss function is thus composed of two different loss functions: (1) $\mathcal{L}_{reconstruct}$ trained on both letters and

⁵<https://emblematic.org/markmaker/>

⁶These font files (TTF) were obtained from a designer colleague.



Figure 4: Example TReATs generated by our approach for (word & theme) pairs: a) (canoe & watersports) b) (world & countries, continents, natural wonders) c) (water & drinks) d) (church & priest, nun, bishop)

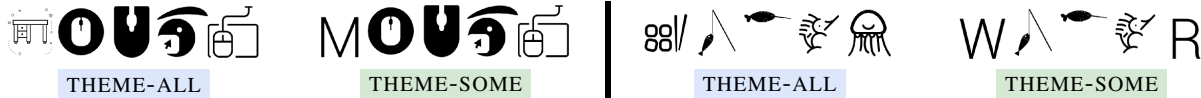


Figure 5: We replace letters in a word with cliparts only if the clipart is sufficiently similar to the letter, placing more stringent conditions on the first and last letters in the word. Notice that in each pair, the TReATs on the right (with a subset of letters replaced) are more legible (Mouse and Water) than the ones on the left, while still depicting the associated themes (computer and fish, mermaid, sailor).

cliparts, and (2) $\mathcal{L}_{\text{classify}}$ trained only on letters. Here α is a tunable hyperparameter in the range $[0, 1]$. We set α to 0.25 after manually inspecting outputs of a few word-theme pairs we used while developing our model (different from the word-theme pairs we use to evaluate our approach later).

Implementation Details: Our encoder network $f_{\text{enc}}(\cdot)$ is an AlexNet (Krizhevsky, Sutskever, and Hinton 2012) convolutional neural network trained from scratch, made up of 5 convolutional and 3 fully connected layers. Our decoder network f_{dec} consists of 5 deconvolutional layers, 3 fully connected layers and 3 upsampling layers. We use batch norm between layers.⁷ We use ReLU activations for both the encoder and decoder. We use the Adam optimizer with a learning rate of 10^{-4} , and a weight decay of 10^{-5} . The input dataset is divided into minibatches of size 100 with a mixture of clipart and letter images in each minibatch. We use early stopping based on a validation set as our stopping criterion.

Data Preprocessing: We resize our images to 224×224 using bilinear interpolation to match the input size of our AlexNet-based encoder. We normalize every channel of our input data to fall in $[-1, 1]$.

Finding Matches

At test time given a word and a theme, we retrieve a theme-relevant pool of cliparts (denoted by \mathbf{C}) by querying Noun Project. These theme-relevant cliparts may or may not be part of the randomly downloaded $\sim 50\text{k}$ cliparts used during training. If multiple phrases have been used to describe a theme, we use each phrase separately as a query. We limit this retrieval to no more than 10,000 cliparts for each phrase. We then combine cliparts for different phrases of a theme together to form the final pool of cliparts for that theme. For example, for the theme `countries, continents, natural wonders`, we query Noun Project for `countries`, `continents` and `natural wonders` individually and combine all retrieved cliparts together to form the final pool of theme-relevant cliparts. On average across 95 themes we experimented with, we had a minimum of 49 and maximum of 29,580 cliparts per theme, with a mean of 9731.2 and median of 9966. We augmented this set of cliparts with left-right (mirror) flips of the cliparts.

⁷Implementations of our encoder and decoder were adapted from <https://github.com/arnaghosh/Auto-Encoder>.

This improves the overall match quality. E.g., in cases where there existed a good match for J , but not for L , the clipart match for J , when flipped, served as a good match for L . Similarly for S and Z .

For each letter l of the input word, we choose the corresponding letter images (denoted by \mathbf{F}_l) taken from a predefined pool of fonts. To create the pool of letter images \mathbf{F}_l , we used uppercase letters from 14 distinct, readable fonts from among the 777 fonts used during training. These were kept fixed for all experiments. We found that uppercase letters had better matches with the cliparts (lower cosine distance between corresponding latent representations Z_i on average, and visually better matches). Moreover, we found that in several fonts, letters were the same for both cases.

We replace the letter with a clipart (chosen from \mathbf{C}) whose mean cosine distance in the intermediate latent space is the least, when computed against every letter image in \mathbf{F}_l . Concretely, if Z_i denotes the intermediate representation for the i^{th} image in \mathbf{F}_l and Z_c is the intermediate representation for a clipart c in \mathbf{C} , the chosen clipart \hat{c}_l is

$$\hat{c}_l = \underset{c \in \mathbf{C}}{\text{argmin}} \frac{1}{|\mathbf{F}_l|} \sum_{i \in \mathbf{F}_l} \left(1 - \frac{Z_i \cdot Z_c}{|Z_i| |Z_c|} \right) \quad (7)$$

We find a clipart that is most similar to the letter on average across fonts to ensure a more robust match than considering a single most similar font. In this way, each letter in the input word is replaced by its closest clipart to generate a TReAT.

We show example TReATs generated by our approach in Figure 4. We find that the word can often be difficult to recognize from the TReAT if the Noun Project cliparts corresponding to a theme are not sufficiently similar to the letters in the word. To improve the legibility of our TReATs, we first normalize the cosine distance values of our matched cliparts for the alphabet for a specific theme in the range $[0, 1]$. We only replace a letter with its clipart match if the normalized cosine distance between the embedding of the letter and clipart is < 0.75 . It is known that the first and last letters of a word play a crucial role in whether humans can recognize the word at a glance. (Rawlinson 2007) So we use a stricter threshold, and replace the first and last letters of a word with a clipart only if the normalized cosine distance between the two is < 0.45 . Example TReATs with all letters replaced and only a subset of letters replaced can be seen in Figure 5. Clearly, the TReATs with a subset of letters

replaced (**THEME-SOME**) are more legible than replacing all letters (**THEME-ALL**), while still depicting the desired theme. We quantitatively evaluate this in the next section.

Evaluation

We evaluate our entire system along three dimensions:

- How well is our model able to learn a representation that captures visual features of the letters?
- How does our chosen source of cliparts (Noun Project) affect the quality of matches?
- How good are our generated TReATs?

Learnt Representation

We use t-Distributed Stochastic Neighbor Embedding (t-SNE) to visualize our learnt latent representations of letters and cliparts. Among letters, we find that our model clusters letters of different fonts together, while distinguishing between visually dissimilar letters. E.g., Figure 7 visualizes in 2D uppercase O, Q, E and F in the 14 fonts used at test time. As expected, O and Q clusters are close, and E and F clusters are close, but both these sets of clusters are apart. Visualizing letters as well as cliparts, Figure 8 shows that our model is able to learn a representation such that visually similar letter-clipart pairs are close in the latent space.

Effect of source of cliparts

Themes which have fewer cliparts, and hence typically lower diversity and coverage across the letters (e.g. *mythical beast* in Figure 9a) have poorer matches as compared to larger, more diverse themes (e.g. *library* in Figure 9b). Indeed, we see that recognizing the word in a TReAT generated from the former theme is significantly harder than for the latter. Note that the diversity of cliparts is more important than the quantity. Fewer but more diverse cliparts can lead to better TReATs than many but less diverse cliparts.

Quality of TReATs

We now evaluate the quality of TReATs generated by our approach. We developed our approach on a few themes (e.g., *education*, *Harry Potter*, *Halloween*, *Olympics*) and associated words (e.g., *exam*, *always*, *witch*, *play*). We select these arbitrarily as diverse and popular domains. To evaluate our approach in the open world, we collected 104 word-theme pairs from subjects on Amazon Mechanical Turk (AMT). We told subjects that given a word and an associated theme, we have a bot that can draw a doodle. We showed subjects a few example TReATs. We asked subjects to give us a word and an associated theme (to be described in 1-5 comma separated phrases) that they would like to see a doodle for. Example (word & theme) pairs from our dataset are (*environment & pollution*, *dirt*, *wastage*), (*border & USA*), (*computer & technology*). We allowed subjects to use multiple phrases to describe the theme to allow for a more diverse set of cliparts to search from when generating the TReAT. We evaluate our TReATs along three dimensions: 1)

Can subjects recognize the word in the TReAT? 2) Can subjects recognize the theme in the TReAT? and 3) Do subjects find the TReAT creative? We conduct independent studies for each of these to eliminate the influence of one on the other.

We compare our approach (**THEME-SOME**) to a version where we replace all letters in the word with cliparts (**THEME-ALL**) to evaluate how replacing a subset of letters affects word recognition (expected to increase) and theme recognition (expected to remain unchanged or even increase because recognizing the word can aid in recognizing the theme), as well as creativity (expected to remain unchanged or even increase because the associated word is more legible as opposed to gibberish). We also compare our approach to an approach that replaces letters with cliparts, but is not constrained by the theme of interest (**NOTHEME-SOME** and **NOTHEME-ALL**). We find the clipart that is closest across all 95⁸ themes in our dataset to replace the letter. This can result in increased word recognition because letters can find a clipart that is more similar (from a larger pool not constrained by the theme), but will result in lower theme recognition accuracy. Note that theme recognition will still likely be higher than chance because the word itself gives cues about the theme. For no-themed clipart, we compare an approach that replaces all letters (i.e., **NOTHEME-ALL**) as well as only a subset of letters (**NOTHEME-SOME**). Finally, as a point of reference, we evaluate a TReAT that simply displays the word in a slightly atypical font (**FONT**). We expect word recognition to be nearly perfect, but theme recognition as well as creativity to be poor. These five different types of TReATs are shown in Figure 6. This gives us a total of 520 TReATs to evaluate (5 types \times 104 word-theme input pairs). No AMT workers were repeated across any of these tasks.

Word recognition: We showed each TReAT to 5 subjects on AMT. 461 unique subjects participated in the word recognition study. They were asked to type out the word they see in the TReAT in free-form text. Notice the open-ended nature of the task. Performance of crowd-workers for word recognition of different types of TReATs is shown in Figure 10a. This checks for exact string matching (case-insensitive) between the word entered by subjects and the true word. As a less stringent evaluation, we also compute individual letter recognition accuracy. These were computed only for cases where the length of the word entered by the subject matched the true length of the TReAT because if the lengths do not match, the worker likely made a mistake or was distracted. Letter recognition accuracies are shown in Figure 10b.

As expected, leaving a subset of the letters unchanged leads to a higher recognition rate for (**THEME-SOME** and **NOTHEME-SOME**) compared to their counterparts, (**THEME-ALL** and **NOTHEME-ALL**) respectively. Also, (**NOTHEME-ALL** and **NOTHEME-SOME**) have higher word recognition accuracy than (**THEME-ALL** and **THEME-SOME**) because the clipart matches are obtained from a larger pool (across all themes rather than from a specific theme). The added signal from the theme of the cliparts in (**THEME-ALL** and **THEME-SOME**) does not help word

⁸95 since some themes repeat in our 104 (word & theme) pairs.



Figure 6: We evaluate five different approaches for generating TReATs. See text for details.




Figure 7: t-SNE plot showing clusters of uppercase O, Q, E and F. Each letter forms its own cluster and visually similar pairs (E & F, O & Q) form super-clusters. However, these super-clusters are far apart from each other due to significant visual differences.

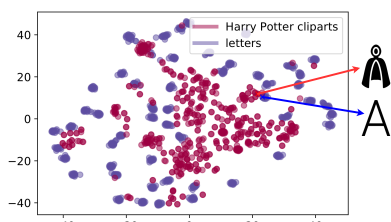


Figure 8: t-SNE plot showing clusters of Harry Potter themed cliparts along with letters. Cliparts which look like A lie close to the cluster of A's in the latent space.

recognition enough to counter this. **NOTHEME-ALL** already has a high recognition rate, leaving little scope for improvement for **NOTHEME-SOME**. Finally, **FONT** has near perfect word recognition accuracy because it contains the word clearly written out. It is not a 100% because of typos on the part of the subjects. In some cases we found that subjects did not read the instructions and wrote out the theme instead of the word itself across all TReATs. These subjects were excluded from our analysis.

Theme recognition: We showed each TReAT to 6 subjects on AMT. 163 unique subjects participated in the theme recognition study. The same theme can be described in many different ways. So unlike word recognition, this task could not be open-ended. For each TReAT, we gave subjects 6 themes as options from which the correct theme is to be identified. These 6 options included the true theme from the 95 themes in our dataset, 2 similar themes, and 3 random themes. The similar themes are the 2 nearest neighbor themes to the true theme in *word2vec* (Mikolov et al. 2013) space. *word2vec* is a popular technique to generate vector representations of a word or “word embeddings” which capture the meaning of the word such that words that share common contexts in language (that is, likely have similar meaning) are located in close proximity to one another in the space. If a theme is described by multiple words, we represent the theme using the average *word2vec* embedding of each word. This is a strategy that is commonly employed in natural language processing to reason about similarities between phrases or even entire sentences (Wieting et al. 2016), (Adi et al. 2017). For example, the options for Figure 4c were 1) home, furniture,

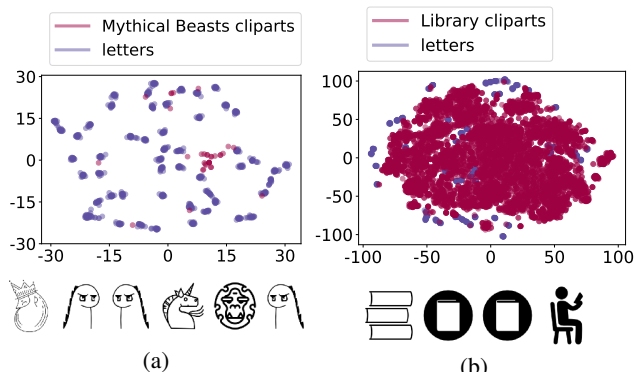


Figure 9: Impact of diversity of cliparts from different themes in the Noun Project on corresponding TReATs. a) TReAT of dragon in theme mythical beast is not legible due to lower coverage of letters by the themed cliparts compared to a TReAT of book in theme library shown in b).

2) drinks, 3) corpse, undertaker, vampire, 4) food, dessert, 5) birds and 6) food with the correct answer being drinks.

We find that 64% of the TReATs were assigned to the correct theme for **THEME-ALL**, 67% for **THEME-SOME**, 43% for **NOTHEME-ALL**, 51% for **NOTHEME-SOME** and 60% for **FONT** respectively. As expected, **NOTHEME-ALL** and **NOTHEME-SOME** have lower theme recognition accuracy than **THEME-ALL** and **THEME-SOME** because **NOTHEME-ALL** and **NOTHEME-SOME** do not use cliparts from specific themes. Notice that theme recognition accuracy is still quite high, because the word itself often gives away cues about the theme (as seen by the theme recognition accuracy of **FONT** that lists the word without any clipart).

This theme recognition rate is a pessimistic estimate because theme options presented to subjects included nearest neighbors to the true theme as distractors. These themes are often synonymous to the true theme. As a less stringent evaluation, we sort the 6 options for each TReAT based on the number of votes the option got across subjects. Figure 10c shows the Mean Reciprocal Rank of the true option in this list (higher is better). We also show Recall@K in Figure 10d that compute how often the true option is in the top-K in this sorted list. Similar trends as described above hold.

Comparing **THEME-SOME** to **THEME-ALL**, we see that replacing only a subset of letters does not hurt theme recognition (in fact, it improves slightly), but improves word recognition significantly. So overall, **THEME-SOME** produces the best TReATs. We see this being played out when TReATs are evaluated for their overall creativity (next). This relates to Schmidhuber’s theory of creativity (Schmidhuber 2010). He argues that data is creative if it exhibits both a learnable or recognizable pattern (and is hence compressible), and novelty. **THEME-SOME** achieves this balance.

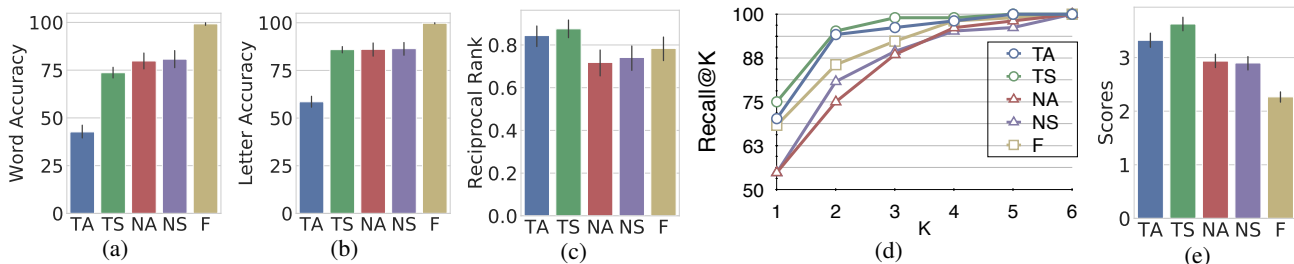


Figure 10: Evaluation of TReATs from five approaches (THEME-ALL (TA), THEME-SOME (TS), NOTHEME-ALL (NA), NOTHEME-SOME (NS), FONT (F)) for a) word recognition; b) letter recognition; c) and d) theme recognition; e) creativity.

Creativity: Recall that our goal here is to create TReATs to depict words with visual elements such that the TReAT leaves an impression on people’s minds. We now attempt to evaluate this. Do subjects find the TReAT intriguing / surprising / fun (i.e., creative)? We showed each TReAT to 5 subjects on AMT. 207 unique subjects participated in the creativity study. They were told: “This is a doodle of [word] in a [theme] theme. On a scale of 1-5, how much do you agree with this statement? This doodle is creative (i.e, surprising and/or intriguing and/or fun). a. Strongly agree (with a grin-like smiley face emoji in green) b. Somewhat agree (with a smiley face in lime green) c. Neutral (with a neutral face in yellow) d. Somewhat disagree (with a slightly frowning face in orange) e. Strongly disagree (with a frowning face in red).” The associated scores were 5 to 1 respectively. Crowd-worker scores are shown in Figure 10e.

THEME-SOME was rated the highest. We believe this is due to a good trade off between legibility and having a theme-relevant depiction that allows for semantic reinforcement. NOTHEME-ALL and NOTHEME-SOME are significantly worse. Recall that they are visual, but not in a theme-specific way. So they are visually interesting, but do not allow for semantic reinforcement. The resultant reduction in creativity is evident. Interestingly, NOTHEME-SOME scores slightly higher than NOTHEME-ALL. This may be because NOTHEME-SOME is not more legible than NOTHEME-ALL (NOTHEME-ALL is already sufficiently legible). With more of the letters visually depicted, NOTHEME-ALL is more interesting. Finally, FONT has a significantly lower creativity score. It is rated lower than neutral, close to the “Somewhat disagree” rating. To get a qualitative sense, we asked subjects to comment on what they think of the TReATs. Some example comments:

THEME-ALL : “cool characters and each one fits the theme of the ocean”, “Its [sic] creative and represents the theme well, but I don’t see disney all that much.”

THEME-SOME : “I like how it uses the image of the US and then a state to spell out the word and looks like something you’d remember.”, “Very fun and intriguing. I like how all the letters are pictures representing a computer mouse.”.

NOTHEME-ALL : “It is creative but it has nothing to do with fear.”, “It does a very good job of spelling out CHRISTMAS, but the individual letters are not related to the holiday at all.”

NOTHEME-SOME : “It is somewhat creative, especially the unicorn head for “G”, though I don’t know what any of it has to do with the theme.”, “There are too many icons that

seemingly have nothing to do with the theme.”

FONT : “It just spells out the word, not really a doodle”, “Its not a doodle, its just the word Parrot, so I don’t think its creative at all.”

Future Work

In this section, we discuss some drawbacks of our current model and potential future work.

No Clipart Relevance Score: A comment from a subject evaluating the creativity of Figure 11a (word church & theme pastor, Jesus, people, steeple) was “you need a cross [...] before the general public would [...] get this.” Our approach does not include world knowledge that indicates which symbols are canonical for themes (Noun Project does not provide a relevance score). As a result, our model can not explicitly trade off visual similarity (VS) for theme relevance (TR) – either to compromise on VS to improve TR, or to at least optimize for TR if VS is poor.

No Contextual Querying: Multiple phrases used to describe a theme often lose context when they are individually queried into the Noun Project. For example, the clipart in Figure 11b for (word money & theme finance, banking, support) contains the image of a cheerleader, which is relevant to the phrase support, but is not relevant in the context of the finance theme.

The lack of context also hurts polysemous theme words. bat when used as a keyword with the another keyword bird refers to the creature bat, but in the context of baseball refers to sports equipment.

Imperfect Match Scores: Our automatic similarity score frequently disagrees with our (human) perceptual notion of similarity. E.g., in Figure 11 right, the cliparts used to replace C and N in THEME-ALL look sufficiently similar to the corresponding letters. But the automatic similarity score was low, and so THEME-SOME chose to not replace the letters. Approaches to improve the automatic score can be explored in future work. For instance, in addition to mirror images, using rotated and scaled versions of the cliparts to augment the dataset would help.

Interactive Interface

To mitigate these concerns, we plan to build an interactive tool. Users can choose from the top- k clipart matches for each letter. Users can iterate on the input theme descriptions until they are satisfied with the TReAT. Users can also leave




Figure 11: Example failure modes of our approach – a) more theme-relevant icons such as the cross should be used to depict the theme *pastor*, *Jesus*, *people*, *steeple*; b) lack of context wherein *support* here actually refers to *financial* support, and not the motivational support which comes from a cheerleader depicting *y* in the word *money*; Right: C and N are replaced in our final **THEME-SOME** TReAT even when the matches were actually quite relevant and visually fitting.

the theme unspecified in which case we can use the word itself as the theme. Finally, users can choose which letters to replace in **THEME-SOME** like TReATs.

Conclusion

In this work, we introduce a computational approach for semantic reinforcement called TReAT – Thematic Reinforcement for Artistic Typography. Given an input word and a theme, our model generates a “doodle” (TReAT) for that word using cliparts associated with that theme. We evaluate our TReATs for word recognition (can a subject recognize the word being depicted?), theme recognition (can a subject recognize what theme is being illustrated in the TReAT?), and creativity (overall, do subjects find the TReATs surprising / intriguing / fun?). We find that subjects can recognize the word in our TReATs 74% of the time, can recognize the theme 67% of the time, and on average “Somewhat agree” that our TReATs are creative.

References

- [Adi et al. 2017] Adi, Y.; Kermany, E.; Belinkov, Y.; Lavi, O.; and Goldberg, Y. 2017. Fine-grained analysis of sentence embeddings using auxiliary prediction tasks. In *ICLR*. 6
- [Atarsaikhan et al. 2017] Atarsaikhan, G.; Iwana, B. K.; Narusawa, A.; Yanai, K.; and Uchida, S. 2017. Neural font style transfer. In *ICDAR*. 2
- [Atarsaikhan, Iwana, and Uchida 2018] Atarsaikhan, G.; Iwana, B. K.; and Uchida, S. 2018. Contained neural style transfer for decorated logo generation. In *IAPR DAS*. 2
- [Azadi et al. 2018] Azadi, S.; Fisher, M.; Kim, V. G.; Wang, Z.; Shechtman, E.; and Darrell, T. 2018. Multi-content gan for few-shot font style transfer. In *CVPR*. 2
- [Ballard 1987] Ballard, D. H. 1987. Modular learning in neural networks. In *AAAI*. 3
- [Campbell and Kautz 2014] Campbell, N. D., and Kautz, J. 2014. Learning a manifold of fonts. *Trans. on Graphics*. 2
- [Clawson et al. 2012] Clawson, T. H.; Leafman, J.; Nehrenz Sr, G. M.; and Kimmer, S. 2012. Using pictograms for communication. *Military medicine*. 2
- [Cunha et al. 2018] Cunha, M.; Martins, P.; João; and Machado, P. 2018. How shell and horn make a unicorn: Experimenting with visual blending in emoji. In *ICCC*. 2
- [Frutiger 1989] Frutiger, A. 1989. Signs and symbols. *Their design and meaning*. 2
- [Ha and Eck 2018] Ha, D., and Eck, D. 2018. A neural representation of sketch drawings. In *ICLR*. 2
- [Isola et al. 2017] Isola, P.; Zhu, J.-Y.; Zhou, T.; and Efros, A. A. 2017. Image-to-image translation with conditional adversarial networks. In *CVPR*. 2
- [Krizhevsky, Sutskever, and Hinton 2012] Krizhevsky, A.; Sutskever, I.; and Hinton, G. E. 2012. Imagenet classification with deep convolutional neural networks. In *NIPS*. 4
- [Martins et al. 2015] Martins, P.; Urbancic, T.; Pollak, S.; Lavrac, N.; and Cardoso, A. 2015. The good, the bad, and the aha! blends. In *ICCC*. 2
- [Martins et al. 2018] Martins, T.; Correia, J.; Costa, E.; and Machado, P. 2018. Evotype: towards the evolution of type stencils. In *International Conference on Computational Intelligence in Music, Sound, Art and Design*. 2
- [Mikolov et al. 2013] Mikolov, T.; Sutskever, I.; Chen, K.; Corrado, G. S.; and Dean, J. 2013. Distributed representations of words and phrases and their compositionality. In *NIPS*. 6
- [Mirza and Osindero 2014] Mirza, M., and Osindero, S. 2014. Conditional generative adversarial nets. *arXiv preprint arXiv:1411.1784*. 2
- [Rawlinson 2007] Rawlinson, G. 2007. The significance of letter position in word recognition. *IEEE Aerospace and Electronic Systems Magazine*. 4
- [Sage et al. 2018] Sage, A.; Agustsson, E.; Timofte, R.; and Van Gool, L. 2018. Logo synthesis and manipulation with clustered generative adversarial networks. In *CVPR*. 2
- [Schmandt-Besserat 2015] Schmandt-Besserat, D. 2015. The evolution of writing. *International Encyclopedia of the Social and Behavioral Sciences: Second Edition*. 2
- [Schmidhuber 2010] Schmidhuber, J. 2010. Formal theory of creativity, fun, and intrinsic motivation (1990–2010). *IEEE Transactions on Autonomous Mental Development*. 6
- [Shiojiri and Nakatani 2013] Shiojiri, M., and Nakatani, Y. 2013. Visual language communication system with multiple pictograms converted from weblog texts. In *IASDR*. 2
- [Takasaki and Mori 2007] Takasaki, T., and Mori, Y. 2007. Design and development of a pictogram communication system for children around the world. In *Intercultural collaboration*. 2
- [Wieting et al. 2016] Wieting, J.; Bansal, M.; Gimpel, K.; and Livescu, K. 2016. Towards universal paraphrastic sentence embeddings. In *ICLR*. 6
- [Xiao and Linkola 2015] Xiao, P., and Linkola, S. 2015. Vis-mantic: Meaning-making with images. In *ICCC*. 2
- [Zhu et al. 2017] Zhu, J.-Y.; Park, T.; Isola, P.; and Efros, A. A. 2017. Unpaired image-to-image translation using cycle-consistent adversarial networks. In *ICCV*. 2