Computational Cut-Ups: The Influence of Dada

By Laure Thompson and David Mimno

Abstract

Can a tool designed to detect dogs detect Dada? We apply a cutting-edge image analysis tool, convolutional neural networks (CNNs), to a collection of page images from modernist journals. This process radically deforms the images, from cultural artifacts into lists of numbers. We determine whether the system can, nevertheless, distinguish Dada from other, non-Dada avant-garde, and in the process learn something about the cohesiveness of Dada as a visual form. We can also analyze the "mistakes" made in classifying Dada to search for the visual influence of Dada as a movement.

> To make a Dadaist poem:
>
> Take a newspaper.
> Take a pair of scissors.
> Choose an article from the newspaper about the same length as you want your poem.
> Cut out the article.
> Then carefully cut out each word of the article and put them in a bag.
> Shake gently.
> Then take out each word, one after another.
> Copy them conscientiously in the order drawn.
> The poem will be like you.
> And look! You are an infinitely original writer with a charming sensibility, yet beyond the understanding of the vulgar.
>
> —Tristan Tzara[1]

Introduction

Can a work of art that has been deformed beyond recognition nevertheless be recognizable? The idea of a cut-up poem is distinctively Dada in its playful reductiveness and, simultaneously, shockingly relevant. Today the cut-up is a foundation of modern textual analysis in the form of the "bag of words" assumption. Search engines, spam filters, and social media recommenders all rely on the assumption that the information carried by the words themselves is sufficient, and that the order in which words appear is irrelevant and burdensome. The bag-of-words assumption reduces the need for intelligence. All that is required is conscientiousness, which computers have in limitless quantities.

In this work we study a deformative "reading" of Dada, not using scissors but a modern computational image-processing method known as a convolutional neural network (CNN). We assure the reader that CNNs are both charming and quite beyond the understanding of the vulgar. CNNs operate by passing images through multiple layers of pattern detectors. The output of a given layer becomes the input of the next layer. For example, the output of the first layer might identify the presence of lines or edges at different angles, while the output of the second layer might identify the presence of pairs of lines that form angles. At the top layer, the output might identify specific things—dog breeds, dishwashers, doormats.[2]

Instead of physical newspapers we cut up digitized avant-garde periodicals from Princeton's Blue Mountain Project.[3] Our initial corpus contains more than 2,500 issues from 36 different journals—over 60,000 pages in total. We deform page-level images into *computational cut-ups* using a CNN. We then use statistical classification to determine which visual features are captured by these cut-ups. Finally, we "read" these computational cut-ups to determine whether such reductive analyses are sufficient to separate Dada from other modernist movements. Our goal is not necessarily to get the "right" answer, but rather to use computation to provide an alien, defamiliarized perspective that can call into question the boundaries between categories.

Creating and Reading Computational Cut-Ups

A computational cut-up is a mathematical representation of an object: a list of numbers that collectively preserve information of the original object. Each value in this list corresponds to a computational feature. In a text model, a feature might correspond to the number of occurrences of a particular word in a document, but image features are more abstract and less apparent. Accordingly, we do not choose these features by hand—we ask a CNN.

CNNs are powerful tools for analyzing images. Although the output of the final layer of a CNN will identify the categories of the object that it was trained to recognize, the output of the next-to-last layer has been shown to produce powerful, high-level visual features. These features are generic enough that they can be used by other image analysis systems.[4] By using these features as our computational cut-ups, we will in essence be asking what CNNs "see" when they look at Dada and more broadly the avant-garde.

Generating a computational cut-up involves two steps. We first shrink the input image to a small 224-by-224 pixel square so that it can be fed as input to our CNN. Fine-grained details are lost through this deformation, but major elements such as layout, headers, and illustrations are generally preserved. As seen in Figure 1, the images fed to the CNN remain recognizable but are similar to viewing the page from the far side of a room.

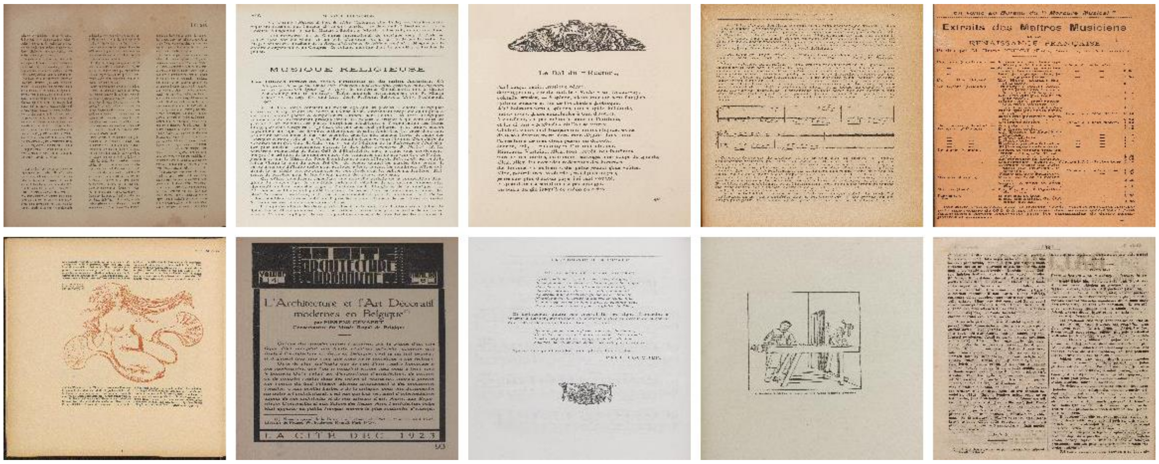

Figure 1: CNN input images for ten randomly sampled pages.

Having created a shrunken version of the original page, we pass it to our CNN as input and extract a computational cut-up from the CNN's internal features.[5] These features are not readily interpretable to the human eye, but they correspond to high-level concepts such as human faces, flowers, and fields of grass.[6] For each feature, we extract a number representing the feature's measured presence within an image—a large value indicates the feature is strongly detected, while a value near zero indicates its absence. As a result, our computational cut-ups are lists of 2,048 numbers.

We evaluate the information encoded within the computational cut-ups by how well they perform on a series of binary classification tasks (e.g. Dada or not-Dada). We measure the degree to which a classifier can distinguish cut-ups with label A from cut-ups with label B. For each computational cut-up *c* with label A or B, we train a Naïve Bayes classifier on all other cut-ups with labels A and B and use it to predict the label for cut-up *c*.[7] The classifier consists of the mean and variance of each of the 2,048 CNN features, for each label. If the feature values of *c* look more similar to the typical feature values of label A, we predict A, and vice versa. The accuracy of these predictions will indicate how well the cut-ups differentiate the two label classes.

In addition to the simple question of whether a classifier is guessing correctly, we are also interested in how confidently the system makes its predictions. We therefore also measure classifier confidence for each prediction. By examining the corresponding page images for the best and worst predictions for each label, we can better understand the visual features being associated with each label.

Proof of Concept: Seeing Music

Before testing whether a CNN can recognize Dada, we verify that it is capable of performing a simpler task: identifying music within periodicals. It is fairly easy for a person to tell the difference between pages of musical scores and pages containing text and images, but how well will our CNN fare? If our computational cut-ups do not distinguish between musical scores and paintings, it would be hard to trust their capability to distinguish Dada from Cubism.

Detecting music within our corpus is a relevant task, not only because music is an avant-garde art form, but because the Blue Mountain corpus has a substantial number of music journals. The five periodicals *La Chronique Musicale*, *Dalibor*, *Le Mecure* (*SIM*), *Niederrheinische Musik-Zeitung*, and *La Revue Musicale* are represented in the corpus by 1,405 issues and 27,791 pages. It is safe to assume that the majority of pages containing music will come from these five journals. Using the TEI-encoded transcriptions for each periodical issue, we identify 3,450 pages containing music.[8] Only 91 of these pages come from the 31 other periodicals.

We find that computational cut-ups are useful for recognizing pages containing music. The classifier makes mistakes that a human might not, but in ways that provide intuition about what it "sees." The classifier correctly labels 67% of the 3,450 pages with music as "Music" and 96% of the 55,007 pages without music "Not-Music." For each prediction, we can measure our classifier's confidence in terms of how much more likely it thinks a page should be labeled as "Music" rather than "Not-Music". Confidence scores with large magnitudes indicate a more confident classification, while a score's sign indicates its assigned label type. So, a large, positive confidence score indicates that the classifier is very confident that a page be labeled "Music." In Figure 2 we see that our classifier tends to be more confident when it labels a page as "Not-Music," even when it is wrong. This difference suggests that the cut-ups may better describe features associated with non-music page elements than music page elements.

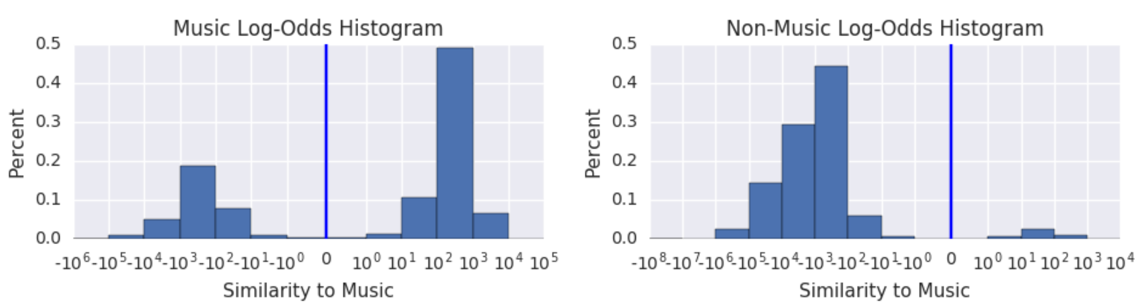

Figure 2: Histograms of prediction confidence for pages containing music (left) and pages without music (right). The classifier is more confident labeling pages as "Not-Music" no matter what the actual page type is.

To understand where the classifier goes wrong we compare the pages that are most confidently classified and misclassified for each label. In Figures 3 and 4, we see that pages of sheet music are most confidently recognized as "Music" and tables are most confidently misclassified as "Music." These images share two prominent features:

prominent horizontal lines and rectangular blank spaces. Given that the actual musical notes are poorly preserved in the deformed CNN inputs, it is reasonable that these are not the dominant visual features associated with pages containing music.

Turning to the “Not-Music” label, we find color and pictures are the dominant visual features associated with pages without music. In Figure 5, we see that the top ten pages most confidently misclassified as “Not-Music” all contain pictures. Moreover, these pictures take up as much space within the page if not more than the musical elements. Many of these pages are also accompanied by text.

Figure 3: Ten pages most confidently, and correctly, classified as “Music.”

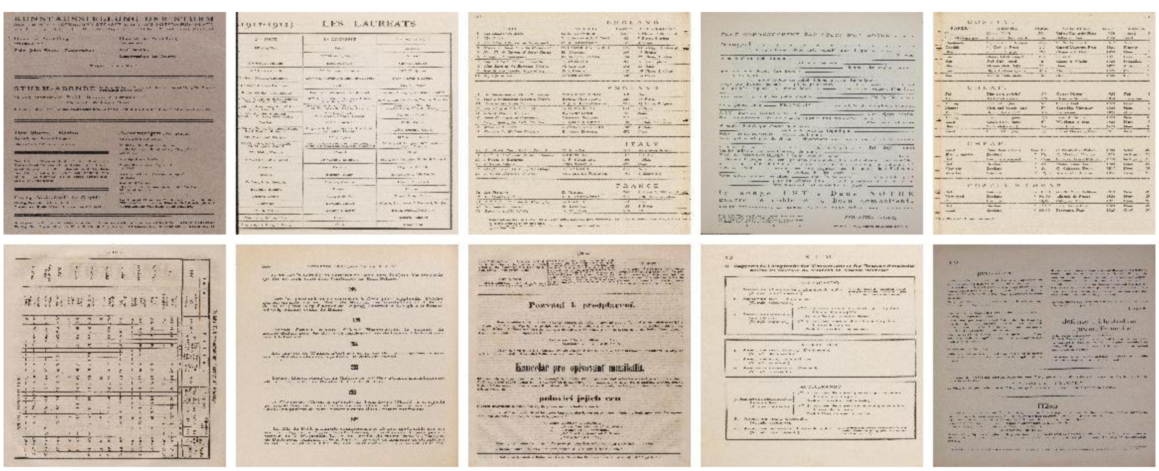

Figure 4: Ten pages most confidently misclassified as “Music.”

Perhaps the most interesting of these confident “Not-Music” misclassifications is the bottom-right page in Figure 5, a scaled down image of a medieval folio. The rescaled CNN input image hardly looks like music, and, in a way, it is not. But looking at the original image in Figure 6, we see it does contain music, even though it looks nothing like modern musical notation. Additionally, the music is being seen through another medium: a picture, which could be misleading the classifier to the “Not-Music” label.

An effective, but potentially misleading feature learned by the classifier is that saturated color indicates no music. While sheet music is generally white, page coloring can vary due to paper and scanning quality. We want to verify that the non-color features perform well without the color cue, and see if the presence of pictures within a page remains the dominant “Not-Music” feature.

We found that removing color from CNN inputs had little effect on classification performance: 66% of pages with music and 97% of pages without music were correctly labeled. Additionally, the most confidently classified and misclassified images remained largely the same for each scenario except for correctly classified pages without music. This is what we had hoped to observe. It indicates that the classifier relies on features other than color. By removing color we also confirmed that the presence of pictures is an important feature for pages without music. As seen in Figure 8, pages containing pictures—both illustrations and photographs—are considered the least musical.

Figure 5: Ten music-containing pages most confidently misclassified as “Not-Music.”

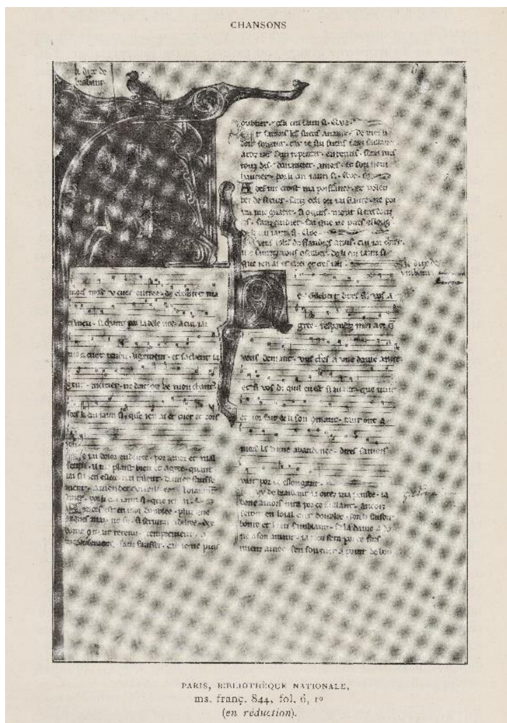

Figure 6: This medieval folio is confidently misclassified as “Not-Music.”

Figure 7: Ten non-music pages most confidently classified as “Not-Music.”

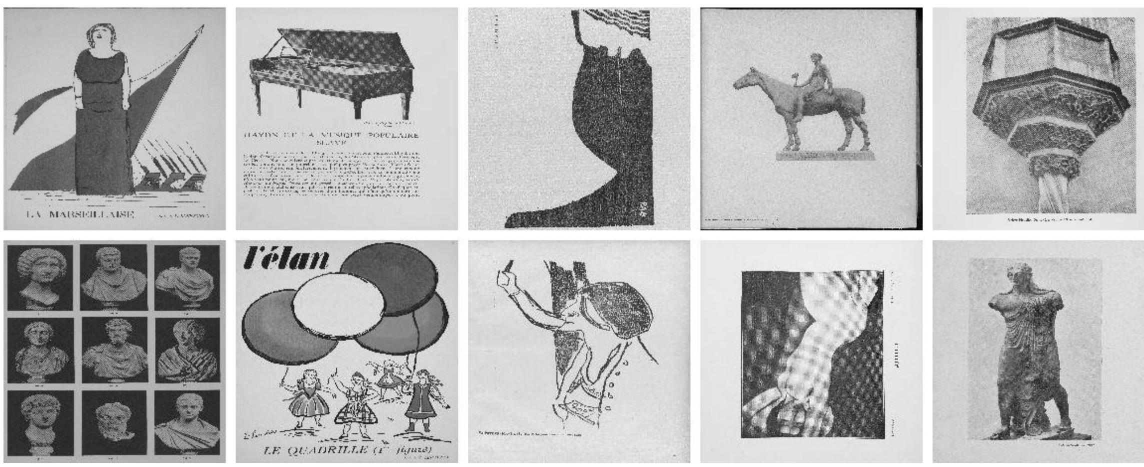

Figure 8: Ten grayscale pages correctly and most confidently classified as “Not-Music.”

Through this exercise, we have shown that computational cut-ups are able to encode visual features that are useful for recognizing pages containing music. Pages with music tend to have regular horizontal lines and rectangular white space, while pages without music tend to contain pictures and be in color. These patterns are fairly primitive, but there is power in this simplicity which echoes the power of word counts to capture abstract textual concepts. Having established our analysis process, we begin our search for what makes Dada Dada.

Distinguishing Dada

For our reading of Dada we begin with the question of whether we can distinguish "Dada" from "Not-Dada." We define labels at the periodical level: for the purposes of this study, the *Dada*, *291*, *Proverbe,* and *Le couer à barbe* are "Dada" and all other periodicals are "Not-Dada". We acknowledge that this is a particularly coarse-grained perspective. A number of periodicals may feature works of Dada artists in specific issues, and these four periodicals might not always feature Dada artists, but these mistakes should have little effect on our classifier given the volume of actual "Not-Dada" material.

We exclude the five music journals from our analysis. Their sheer volume in the Blue Mountain Project would likely drown out the visual features that we are most interested in finding. Moreover, we would like to avoid learning the naïve feature that Dada does not contain sheet music, and hopefully uncover more interesting distinctive features. After this exclusion, we have 32,642 pages labeled "Not-Dada" and 182 pages labeled "Dada."

We find that computational cut-ups are not perfect at distinguishing "Dada" from "Not-Dada," but they are better than random. The classifier correctly labels 63% of the Dada pages and 86% of the not-Dada pages. In Figure 9, we see that the classifier is, as with music, more confident about its "Not-Dada" predictions. We speculate that other avant-garde movements may have visual signals that are easier to identify than Dada.

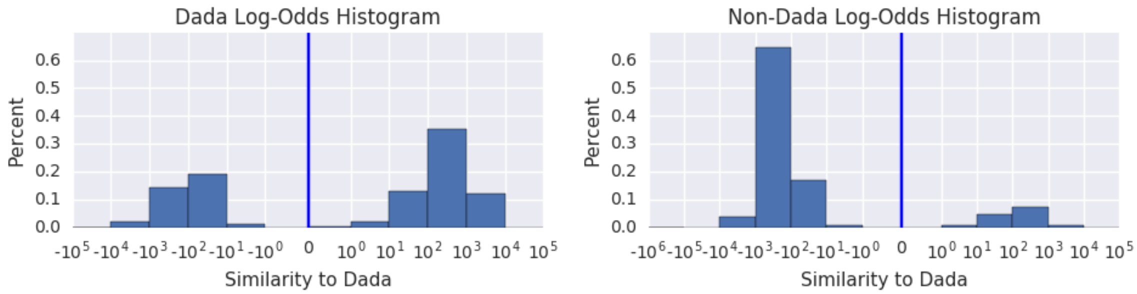

Figure 9: Histograms of prediction confidence for Dada (left) and not-Dada (right) pages. The classifier is more confident labeling pages as “Not Dada” no matter what the actual page type is.

What then does the classifier “see”? When examining the classifier’s most confident successes and mistakes in Figures 10–13, we find that the low-level features associated with Dada are high contrast, prominent edges, and the color red. In comparison, graded texture and photographs are considered not-Dada. From these low-level features, we see that abstract human forms are generally associated with Dada, while more realistic forms are not.

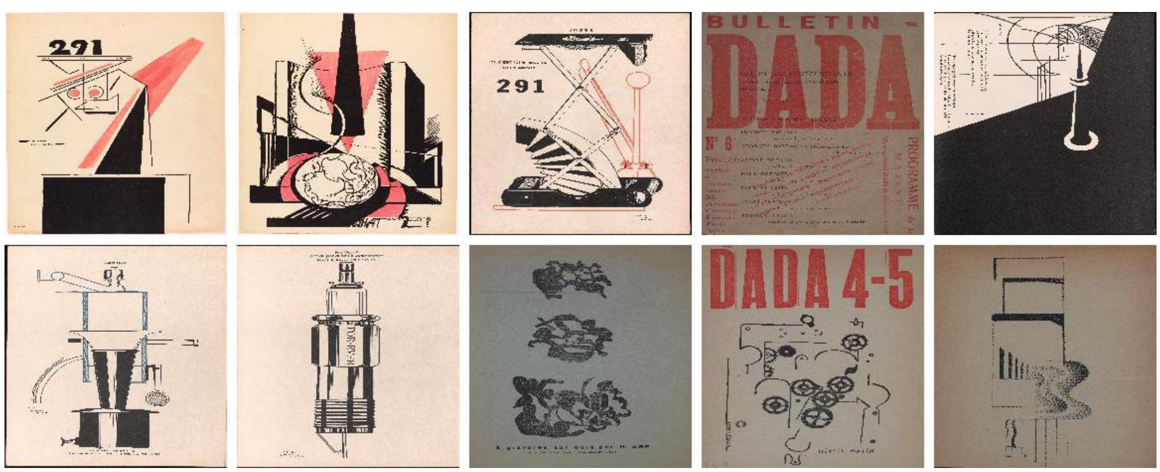

Figure 10: Ten Dada pages most confidently classified as “Dada.”

Figure 11: Top 150 not-Dada pages most confidently misclassified as “Dada.”

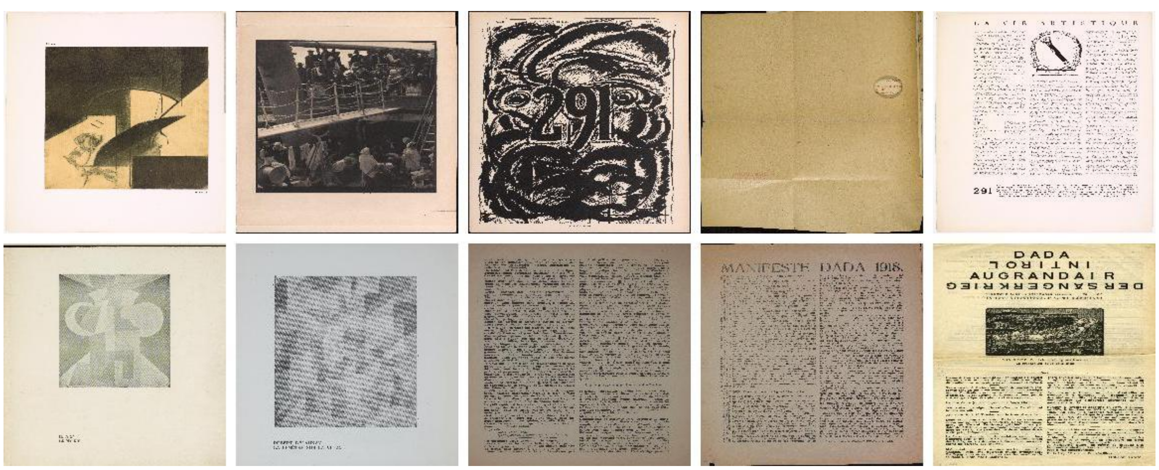

Figure 12: Ten Dada pages most confidently misclassified as “Not-Dada.”

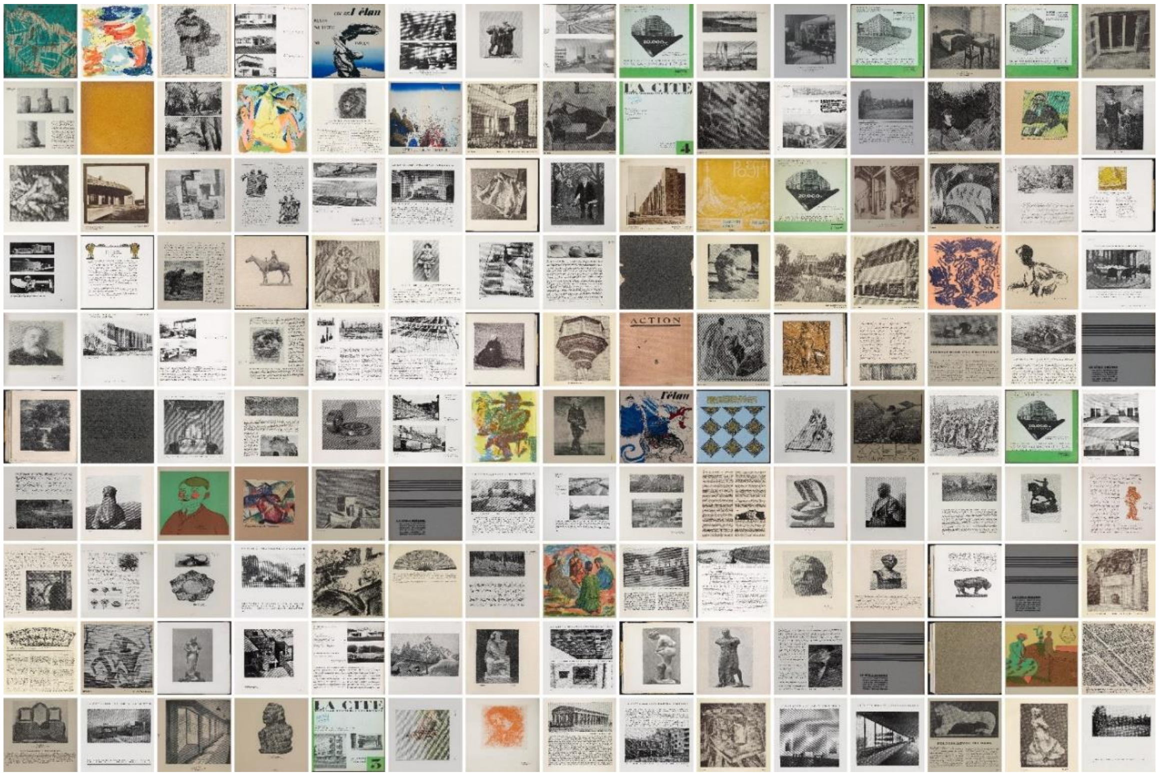

Figure 13: Top 150 not-Dada pages most confidently classified as “Not-Dada.”

Given the prominence of red in “Dada” labeled pages, we were concerned that our results were overly dependent on this simple variable, and not able to generalize to shape or texture. We therefore reran the same analysis on grayscale images to measure the overall effect of color. The classifier’s accuracy worsens for both label groups with resulting accuracies of 56% for “Dada” and 84% for “Not-Dada.” Since this degradation is relatively small, we conclude that color is an important feature for distinguishing Dada, but not the only feature. We find that contrast, edge sharpness, and texture all remain prominent features for classification in grayscale.

It is perhaps unsurprising that color would play a role in distinguishing periodical groups, since page color is influenced by both content and printing method. If a journal has a distinctive page coloring, then it can easily be distinguished from other periodicals by this color alone. This feature can both cause pages with ambiguous content to be correctly identified and pages with otherwise highly similar content to be easily distinguished because of differences in color palette.

We find that our classifier is most confident when labeling pages containing images, and least effective for text-only pages. This result suggests that the classifier is less certain about how text-only pages relate to Dada. Although the CNN does not appear to be able to distinguish between journals based on pages of text, we should not conclude that there are not typographic or layout features that could distinguish them, since these features may simply not be preserved in downscaled pages. We will keep this concept in mind as we analyze our results at the periodical series level.

With these intuitions about which features appear Dadaesque according to the CNN's deformative viewing, we can measure Dada at level of an entire periodical series. Which periodicals "fool" the classifier, and therefore question the (somewhat arbitrary) boundaries that we have constructed? We measure a periodical's closeness to Dada by the proportion of its pages classified as "Dada." As seen in Table 1, we find that the journals *Le coeur à barbe*, *Dada*, *291*, and *L'élan* are the most Dadaesque journals with over half of each periodical's pages classified as "Dada" for both color and grayscale images. Notably, *L'élan* is not a Dada periodical; it is a Cubist war journal. In fact, we find a number of Cubism and Cubism influenced journals among the most Dada-like not Dada, namely *Klingen*, *Der Sturm*, and *SIC*. We find this to be a positive result given Cubism's influence on Dada, particularly Dada art.

We were surprised to find that the two Dada-related journals *Secession* and *Nord-Sud*, as well as the Surrealist journal *Surréalisme* were found to be not Dada-like. While

Nord-Sud's Dada proportion improves with grayscaling, it is still far below fifty percent. Examining these journals, we find that they are predominantly composed of text. Given our suspicions that our classifier has difficulty correctly classifying text-only content, we narrow our measurements to pages containing images.

Table 1: Periodical-level Dada-like proportions for all pages.

Periodical	**Total Pages**	**% Dada (Color)**	**% Dada (Grayscale)**
Le coeur à barbe	8	100.00	75.00
Dada	110	70.91	60.91
291	42	59.52	57.14
Proverbe	22	18.18	18.18
L'élan	129	58.91	56.59
Klingen	727	26.55	31.09
Veshch Gegenstand Objet	64	25.00	29.69
Der Sturm	4649	22.54	24.71
La cité	4766	18.90	21.74
SIC	458	16.16	19.43
Ver sacrum	1928	16.08	21.27
Umělecký měsíčník	933	15.01	18.44
Bruno's Weekly	1219	14.68	15.42
Volné směry	1797	14.47	17.25
Zeit-Echo	756	12.70	17.06
New Numbers	222	11.71	10.36
Poesia	1603	9.61	10.61
Action	911	9.55	14.16
Broom	1751	9.42	10.34
Sturm-Bühne	32	9.38	21.88
Entretiens politiques & littéraires	2764	9.37	8.54
Ultra	156	7.69	11.54
The Mask	3980	7.66	9.92
Nord-Sud	246	7.32	14.63
The Glebe	940	7.02	9.68
Surréalisme	16	6.25	6.25
Secession	217	5.07	5.53
Nowa sztuka	76	3.95	10.53
The Signature	96	3.13	9.38
Pan	2136	3.09	8.19
East & West	70	1.43	0.00

We identify pages containing images using the periodical TEI transcripts similar to our identification of pages containing music. If a page contains content marked as an illustration, then we consider it an "illustrated page."[9] We note that these annotations make different judgements on what content constitutes illustrated content and, as a result, smaller, more decorative illustrations are included inconsistently. Nevertheless, this labeling is sufficient to allow us to demonstrate patterns; future work could specifically analyze the images within journals. We find around one-third of the pages are illustrated, with a very uneven distribution across periodicals.

As we narrow our focus to "illustrated" pages, the Dada-like page proportions increase across journals as shown in Table 2. Encouragingly, the Dada-like journals found in the overall page set remain Dada-like. Now, *Surréalisme*, *Secession*, and *Nord-Sud* have much higher Dada percentages, although these three journals are represented by very few pages. Moreover, the more illustrated journals *Der Sturm* and *SIC* have a high Dada page proportion.

Table 2: Periodical-level Dada-like proportions for "illustrated" pages.

Periodical	**"Illustrated" Pages**	**% Dada (Color)**	**% Dada (Grayscale)**
Le coeur à barbe	1	100.00	100.00
Dada	61	85.25	85.25
291	31	74.19	67.74
Proverbe	3	33.33	33.33
Surréalisme	1	100.00	100.00
Secession	7	100.00	85.71
Der Sturm	1034	74.76	79.40
SIC	83	71.08	78.31
L'élan	87	66.67	73.56
Broom	238	52.94	60.92
Nord-Sud	4	50.00	100.00
Poesia	136	47.06	46.32
Bruno's Weekly	317	45.11	44.79
Ultra	23	39.13	56.52

Zeit-Echo	213	38.97	47.89
Klingen	464	36.85	44.83
La cité	1784	32.23	40.30
Veshch Gegenstand Objet	23	30.43	34.78
Umělecký měsíčník	395	27.85	34.49
Action	222	25.23	40.99
Ver sacrum	1415	16.11	22.83
The Mask	1500	14.20	17.73
Volné směry	1286	12.60	18.04
Pan	1041	5.76	15.37
East & West	1	0.00	0.00
Entretiens politiques & littéraires	1	0.00	0.00
Nowa sztuka	1	0.00	0.00
The Glebe	0	N/A	N/A
New Numbers	0	N/A	N/A
The Signature	0	N/A	N/A
Sturm-Bühne	0	N/A	N/A

Conclusion

A deformative technique such as the cut-up poem seeks meaning in the visible features of language, while playfully ignoring the original concepts and intentions of the text. By "reading" data art with Convolutional Neural Networks (CNNs), we can take the same approach, isolating ourselves from concepts and intentions, and accessing only visual features. The CNN can only describe and distinguish, not define. We find that CNNs indeed enable a deformative viewing of modernist journals. This in itself is not surprising: a tool that analyzes images analyzed images. What is critical from a scholarly perspective is whether this deformative reading provides a perspective that is both distinct from and complementary to human reading. Can a tool designed for identifying dogs be repurposed for exploring the avant-garde? Can it see Dada among the rest?

From the perspective of a computational cut-up of Dada journal pages, we find that pages of Dada journals can be distinguished from pages of non-Dada journals with a degree of accuracy that exceeds random chance. This suggests that there is substance behind the name of Dada. The internal state of neural networks is notoriously inscrutable, but by sorting pages by predicted Dada-ness, we can start to infer how the machine "sees" Dada both from its successes and its mistakes. Dada is characterized (*not* defined) by red hues, sharp and prominent edges, and high contrast. These features are simplistic but can be combined to form more complex structures such as schematic-like figures and abstract human forms. The pages that the CNN thinks really ought to be Dada show the porous boundaries of the category: Cubism thus appears to be measurably the closest movement to Dada at the level of simple visual features.

This characterization of Dada is both alien and familiar. It is produced by an alien gaze, by a machine trained to identify image content ranging from specific dog breeds to microwaves and guillotines. From this machine-view, we gain abstract, but at times unfamiliar, features that nonetheless reflect human concepts. In the case of Dada, the CNN directs our attention to the presence of abstract forms and schematic drawings, and strongly away from photography and more realistic representations of the body. Is this a machine's way of separating art from anti-art?

A potential shortcoming or strength of this machine reading is its illiteracy. The CNN was not trained to read human language; moreover downscaling images makes text largely illegible if not invisible. This prevents the CNN's ability to "cheat" by associating Dada with the name of the movement or artists associated with it. Instead, it must find visual cues that are significant to Dada journals alone. In all likelihood, this causes the CNN to fixate on particular artists and their styles. Clearly, it will take art at face value and not read into the intent of the artist. However, this deconstruction of art echoes the effects of Dadaism itself.

As with all scholarship, but particularly data-driven scholarship, our analysis is limited by the scope of the collection. This reading primarily focused on the pictures contained within a page, and on the journals present. The former is a shortcoming of the CNN, while the latter is of the data itself. The CNN is unlikely to associate the poems of Tristan Tzara or the readymades of Marcel Duchamp as Dada because its attention is focused away from texts and photographs. Similarly, it cannot associate other avant-garde art with Dada that it never sees. Despite these shortcomings, we come upon an interesting and believable finding. From the perspective of the CNN and this collection, Dada is most similar to Cubism. Unfortunately, potential connections to Surrealism could not be observed because the Surrealist journals we included had no pictures.

The idea of viewing art with computers necessarily implies a reductive and even ludic perspective. The choice of Dada as a testbed for this approach is quite deliberate, and one we hope fits with the spirit of the movement. The characteristics of Dada learned by the CNN may simply be artifacts of printing choices, and almost certainly "miss the point" at a conceptual level. But they also force us to recognize the visible, structural characteristics of Dada art, and more importantly, point us to the potential connections and influences of the movement outside Dada proper. The CNN-based classifier is like Dada, but has its own sensibility. Perhaps in explaining its successes and puzzling over its mistakes we may ourselves become infinitely original.

Notes

1. Tristan Tzara. "Pour faire un poème dadaïste." In *Lampisteries, précédées des Sept manifestes Dada: quelques dessins de Francis Picabia* ([Paris]: J. J. Pauvert, 1963): 64-5.

2. See http://image-net.org/challenges/LSVRC/2014/browse-synsets for an example list of object classes used in object detection and image classification tasks.
3. The transcripts contain human-generated metadata which describes the editorial content within a periodical and their corresponding page locations. See http://bluemountain.princeton.edu.
4. Sharif Razavian, Ali, Hossein Azizpour, Josephine Sullivan, and Stefan Carlsson. "CNN features off-the-shelf: an astounding baseline for recognition." In *Proceedings of the IEEE conference on computer vision and pattern recognition workshops* (2014): 806-813.
5. We use the ResNet50 model pre-trained on ImageNet which is available through Keras.
6. See http://yosinski.com/deepvis#toolbox for more information on visualizing neural network features.
7. Peter M. Broadwell, David Mimno and Timothy R. Tangherlini. "The Tell-Tale Hat: Surfacing the Uncertainty in Folklore Classification." *Cultural Analytics* (2017).
8. We consider content marked as "Music" to represent musical content within a page. See https://github.com/cwulfman/bluemountain-transcriptions.
9. We consider content marked as "Illustration" or "ComplexIllustration" to represent images within a page.