

# Classifying Pro-Eating Disorder Images on Twitter Using Deep Learning

Jonathan Feldman  
*Paschal High School*

## Abstract

The Covid-19 pandemic induced a vast increase in adolescents diagnosed with eating disorders and hospitalized due to eating disorders. This immense growth stemmed partially from the stress of the pandemic but also from increased exposure to content that promotes eating disorders via social media, which, within the last decade, has become plagued by pro-eating disorder content. This study aimed to create a deep learning model capable of determining whether a given social media post promotes eating disorders based solely on image data. Tweets from hashtags that have been documented to promote eating disorders along with Tweets from unrelated hashtags were collected. After preprocessing, these images were labeled as either pro-eating disorder or not based on which Twitter hashtag they were scraped from. Several deep-learning models were trained on the scraped dataset and were evaluated based on their accuracy, F1 score, precision, and recall. Ultimately, the Vision Transformer model was determined to be the most accurate, attaining an F1 score of 0.877 and an accuracy of 86.7% on the test set. The model, which was applied to unlabeled Twitter image data scraped from "#selfie", uncovered seasonal fluctuations in the relative abundance of pro-eating disorder content, which reached its peak in the summertime. These fluctuations correspond not only to the seasons but also to stressors, such as the Covid-19 pandemic. Moreover, the Twitter image data indicated that the relative amount of pro-eating disorder content has been steadily rising over the last five years and is likely to continue increasing in the future.

## 1 Introduction

### 1.1 Background

The Covid-19 pandemic brought about a significant increase in the number of adolescents diagnosed with eating disorders and in the number of adolescents hospitalized due to eating disorders (Agostino et al., 2021; Matthews et al., 2021; Otto et al., 2021). This growth stemmed from the fact that the pandemic expanded the prevalence of factors that induce eating disorders, such as seclusion and stress, and decreased the availability of factors that would inhibit the development of eating disorders (Culbert et al., 2015). Moreover, the pandemic caused a near-universal growth in social media use (Parsons et al., 2021; Rosen et al., 2022), and since social media serves as an environment that not only fosters pro-eating disorder (Pro-ED) sentiment, but also spreads it worldwide (Branley & Covey, 2017; Sukunesan et al., 2021), Pro-ED content was able to reach more people than ever before. In turn, a greater number of people were exposed to these dangerous ideas and began to post similar content (Mento et al., 2021), growing the Pro-ED movement substantially.

Recently, due to the pandemic induced rise in eating disorder cases, numerous studies were conducted analyzing a vast array of possible methodologies to mitigate the spread of Pro-ED content on social media. Among these studies, several have examined the possibility of applying artificial intelligence-based techniques to mitigate the spread of content that promotes eating disorders on social media (Fardouly et al., 2022; Ren et al., 2022). Many have already applied models such as the Bidirectional Encoder Representations from Transformers (BERT) model and its derivatives to determine whether text-based social media posts are promoting eating disorders (Benítez-Andrades et al., 2021, 2022; Oksanen et al., 2015). Despite the numerous studies of eating disorder related content on social media and the vast improvements in image-based artificial intelligence algorithms, none have yet employed computer vision techniques to detect whether an image in a social media post promotes eating disorders. Images constitute a substantial portion of posts on social networking sites, thus making image analysis crucial to understanding the role of eating disorders online. In addition, images, specifically those with the intent to promote eating disorders, are capable of inducing body dissatisfaction, emotional distress, and detrimental eating habits in viewers, which can be injurious to those suffering from an eating disorder or predisposed to an eating disorder (Spettigue & Henderson, 2004). Hence, it is critical that image-based artificial intelligence techniques be utilized to attenuate the influence of the pro eating disorder movement on social media.

### 1.2 Objectives

This study had two objectives. The first objective was to create an accurate image classifier that could detect content that promotes eating

disorders on social media solely from image data. The second objective was to use the trained classifier to analyze the content of tweets over time to understand to what extent the ubiquity of pro-eating disorder content has grown on twitter over the last five years and to uncover any underlying patterns within Pro-ED movement on twitter.

### 1.3 Literature Review

Supervised machine learning (ML) and deep learning (DL) methods, which can be used to predict an outcome based on input data, have flourished thanks to the availability of data on social media. There are many supervised learning techniques, such as regression and classification (Nasteski, 2017). All of them require a labeled dataset which is used to train, and, possibly, test the model. Supervised learning models can then be used to classify large swathes of data, saving time, and generating predictions of greater accuracy than humans could (Goh et al., 2020).

A previous study (Benítez-Andrades et al., 2022) created a dataset of over two-thousand classified tweets which were subsequently used to train and test several ML and DL models, all of which seemed to be highly accurate. Another study researched the potential benefits and limitations of artificial intelligence for prevention and diagnosis of eating disorders, concluding that such models are highly effective (Fardouly et al., 2022). Several other studies have also used social media data to analyze the effect of eating disorder related content from a social perspective instead of a quantitative one (Gordon, 2000; Khosravi, 2020; Parsons et al., 2021).

In short, previous studies have been successful in detecting posts that promote eating disorders (Pro-ED) on social media using natural language processing (NLP) techniques, showing that DL and ML techniques can accurately recognize Pro-ED text (Benítez-Andrades et al., 2021, 2022). This study differs from its predecessors because, rather than analyzing text posted on Pro-ED platforms, it analyzes images. On social media, images play a vital role in communication and the diffusion of ideas, therefore, image analysis offers crucial insights into the dispersal of Pro-ED content online.

## 2 Methods

### 2.1 Data Collection

Publicly available tweets from several Twitter hashtags were scraped, downloading not only the content of the tweet, but also information about the post, such as the tweet id, the number of likes, retweets, and replies, and hyperlinks to any images associated with the post. This study uses Twitter to obtain the dataset because, unlike other social media platforms such as Instagram, Facebook, and Tumblr, it does not currently have a policy for censoring hashtags or posts that promote eating disorders. Twitter is, therefore, an ideal space for Pro-ED content to disseminate and is host to a massive Pro-ED community (Arseniev-Koehler et al., 2016; Bert et al., 2016; Sukunesan et al., 2021). Such an unregulated platform

provides space for a wholistic analysis of unrestricted Pro-ED content, which is ideal for this study.

Tweets from hashtags that previous studies have shown to disseminate Pro-ED content were scraped in their entirety. These hashtags consisted of “#proana”, “#thinspo”, “#thinspiration”, “#fitspiration”, and “#fitspo”. These specific hashtags were chosen because previous research has revealed them to spread Pro-ED content (Dignard & Jarry, 2021; Sukunesan et al., 2021). Images from these hashtags made up the “Pro-ED” section of the scraped dataset. Similarly, Tweets from hashtags not associated with eating disorders were also scraped. These Tweets encompass many diverse images including those of people, animals, and landscapes. Such hashtags included “#ootd”, “#fakecandid”, “#animals”, “#pets”, “#travel”, and “#photography”. These hashtags have not been shown to spread Pro-ED content and are not associated with the hashtags from the Pro-ED portion of the dataset. These are valid candidates for the Not-Pro-ED section of the dataset because they contain a wide variety of image subjects, which trains the model on many different image types, subsequently decreasing the chances that a model will misclassify an image.

Upon scraping the tweet metadata, the Twitter image links were extracted, and all images were downloaded and marked as Pro-ED or Not Pro-ED based on the hashtags from which they were scraped, following the categorization described earlier. Such an approach, while potentially open to misclassifications, is an efficient way to create large image datasets on which to train the classifier, which is of particular importance when analyzing the manifold content of social media.

## 2.2 Data Cleaning

In total, 18,822 images were scraped from Twitter. Since most of the images were associated with more than one of the scraped hashtags all images with identical image URLs were treated as duplicates and deleted from the dataset. To ensure all duplicates were removed, a dHash algorithm (Yang, 21) was used to remove all images that had a similarity of 90% or more (Chamoso, 2017). Following the removal of duplicate images, the combined dataset decreased to 13,448 images.

## 2.3 Data Labeling

The dataset was subsequently split into two distinct categories based on their associated hashtag. Images from hashtags that are documented to proliferate Pro-ED content, such as “#proana”, were labeled as “Pro-ED”, while images from Not Pro-ED hashtags were labeled as “Not Pro-ED”. Images classified as “Pro-ED” were then relabeled as 0, while images classified as “Not Pro-ED” were relabeled as 1—a necessary step, since classification algorithms can only process numerical values.

### 3 Classification Methods

#### 3.1 General

Two computer vision frameworks were chosen for this classification task, the Residual Neural Network (ResNet) model and the Vision Transformer (ViT) model. The ResNet algorithm was chosen because it has been the default Convolutional Neural Network (CNN) framework since it won the ImageNet Large Scale Visual Recognition Challenge in 2015. The ResNet algorithm also serves as a baseline computer vision model to which other image recognition algorithms can be compared since it has been heavily studied and documented (He et al., 2015; Li et al., 2017; Targ et al., 2016; Wu et al., 2019). The ViT algorithm was chosen because it utilizes a transformer-based framework (Dosovitskiy et al., 2021), unlike the standard CNN, which reflects the current state-of-the-art in image recognition (Yu et al., 2022). Both frameworks were initially trained on the ImageNet-1k dataset (Deng et al., n.d.), which consists of over 1.2 million images. To classify Tweets, each model was individually fine-tuned on the training dataset. Fine-tuning involves freezing all layers of the neural network except the final fully connected layer, and then changing only the weights of the non-frozen final layer to create a classifier suitable for the necessary task. Fine-tuning allows the classifier to utilize pretrained weights from a large dataset such as ImageNet-1k, which aids in recognizing features in images, leading to a greater classification accuracy (Tan et al., 2018; Zhou et al., 2017). Both models were trained on the same combined dataset, of which 80% of the images were used for training and validation, while 20% of the images were used for testing. Each model was fine-tuned for 20 epochs. The weights from the epoch that achieved the highest validation accuracy were used in the final model.

#### 3.2 ResNet

The ResNet model (Rao et al., 2021) follows the classic CNN structure which is characterized by adjoining groups of convolutional layers attached together by a pooling layer. Most CNNs end with a group of fully connected linear layers that output the final classification (He et al., 2015). The basic structure of ResNet was not a new development. It was first seen in the LeNet-5 image classifier which was originally published in 1998. What differentiates ResNet from its predecessors is its approach to mitigating the vanishing gradient problem—a common issue in deep neural networks since the gradient continually decreases as it propagates backwards through the model until it all but vanishes. The ResNet architecture alleviates vanishing gradient problem by providing gradients a means by which they can skip over layers. The gradient then benefits by maintaining its value for longer. For this task, a variant of ResNet, ResNet-152, was used. ResNet-152 is already pretrained and is available through the PyTorch model library. This model differs from other ResNet models due to its increased number of convolutional/pooling layer groups.

As a result of the depth of ResNet-152, it outperforms most ResNet variants (Han et al., 2018), while maintaining the integrity of its gradient due to the skip connections.

### 3.3 Transformer

Transformers—an integral part of Vision Transformer models—are a deep learning model that uses a self-attention mechanism to impute significance to any given section of input data (Vaswani et al., 2017). Self-attention allows transformers to draw information from pieces of input data that came prior, providing crucial knowledge of the relationship between different inputs within the sequence. Transformers analyze sequential data but process the entire input all at once. During processing the self-attention mechanism establishes contextual links between the different parts of the input data, such as words in an input statement. However, due to this self-attention processing, transformers can process multiple input pieces at a time, allowing for parallelization and reducing time spent on computation (Rao et al., 2021).

### 3.4 Vision Transformer

The Vision Transformer (ViT) model differs greatly from the past CNN structures because it employs transformers rather than a sequence of convolutional and pooling layers. ViT, inspired by transformer-based NLP models (Dosovitskiy et al., 2021), leverages a transformer-based architecture for computer vision tasks, achieving a far better performance than regular CNNs. The self-attention mechanism within transformers allows ViT to avoid the common inductive biases of standard CNNs by simultaneously analyzing the entire input image. In essence, transformer-based computer vision models can extract not only visual features from an image, but also interdependencies between different parts of an image, by analyzing the image wholesomely (Islam, 2022; Jeeveswaran et al., 2022).

Unlike ResNet, ViT does not follow the standard CNN structure. Alternatively, the model begins by splitting up an input image into flattened patches and transforming the patches into sequential linear embeddings. These embeddings are then inputted into the transformer encoder which alters the data based on learned interdependencies. Afterwards, the encoded data is fed through a multilayer perceptron (MLP) which generates a classification.

### 3.5 Procedure

The image dataset was organized, loaded, and transformed through PyTorch's built-in imagefolder class. The ResNet-152 model was instantiated using PyTorch's torchvision.models' package. The ViT Model was downloaded from the Google AI's Hugging Face model repository. The ViT model was finetuned through the Hugging Face API in addition to PyTorch's computer vision library: Torchvision.

## 4 Results

### 4.1 Performance

The fine-tuned ViT model was found to be more accurate on the test set than the ResNet model. ViT achieved approximately 86.7% accuracy and received a F1 score of 0.877. The ResNet model achieved 83.2% accuracy and received an F1 score of 0.837. The difference in accuracy and F1 score between the two models is expected since research has indicated that ViT models consistently outperform ResNet models (Chen et al., 2022). Moreover, ViT also routinely attained faster training and inference speeds than ResNet-152, as can be seen in the Error vs. Epoch graphs in Table 1 which graphs the models' accuracy, relative to the correct prediction, over the period of training time, otherwise known as epochs. Such graphs illustrate the gradual improvement of a model's accuracy as it "learns" from the data.

Both the ResNet-152 and the ViT model have a significantly higher recall than precision, indicating that both models are more likely to generate false positive classifications than false negative classifications.

### 4.2 Analysis

After concluding that the ViT model outperforms the ResNet model, a method to test the efficacy of the model on large social media datasets was devised. To examine the ViT model's capabilities, twitter images from "#selfie" were scraped. It was determined that stratified random sampling would be the most effective sampling method, because precise conclusions can be made about certain subgroups, in this case months, from stratified sampling without the need for large datasets (Howell et al., 2020; Shadish et al., 2002). For a given month, three random non-consecutive days were chosen, and for each of the three chosen days, all images posted on the chosen hashtags were scraped.

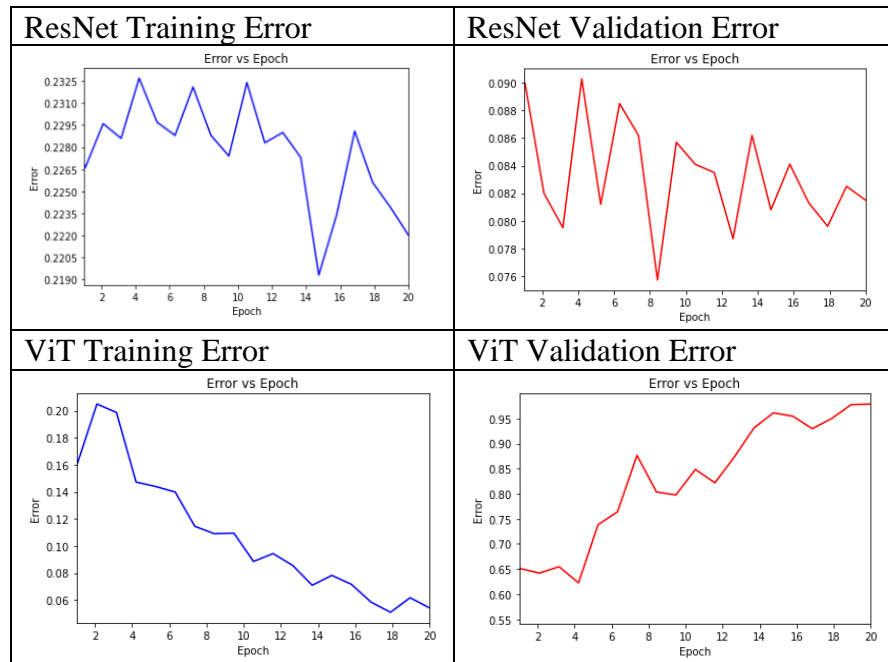


TABLE 1. Error vs. epoch graphs. The blue colored graphs measure the training error of each model, while the red colored graphs measure the validation error of each model.

The process was repeated for each month between January 2017 and June 2022. The hashtag, “#selfie”, was chosen because most of the images posted on it are focused on human beings and their appearance, mimicking the content posted on Pro-ED hashtags. In total, 56,445 images were scraped for analysis from “#selfie”. All the scraped images were then aggregated by month and then fed into the ViT model, which generated a classification. For each month, the classifications for the individual images were used to calculate the percentage of Pro-ED images posted within the time frame. These percentages were then arranged by date and graphed.






FIGURE 1. Linear regression data fit




FIGURE 2. Polynomial regression data fit

From this graphed data, a trend line was generated, shown in Figure 1. The trend line indicates that there is a positive correlation between the progression of time and the relative amount of Pro-ED images on “#selfie”. However, the r-squared value of the trend line is 0.166. Such a small r-squared value reveals that the trend line does account for a vast majority of the variation in the data. To remedy this, a polynomial regression was fit to the data, as shown in Figure 2. The generated regression had an r-squared value of 0.555, indicating that the polynomial graph illustrates most of the data’s variations. Moreover, the RMSE of the linear regression is 1.31, while the RMSE of the polynomial regression is 0.871, once again indicating that a polynomial regression provides a superior fit. Both the linear regression and the polynomial regression have holistic P values of less than 0.001 ( $P < .001$ ), which indicates that both regression models are statistically significant.




FIGURE 3. Aggregate mean of pro-ED content per month from 2017-2022. This figure is the graph of the mean percent of Pro-ED content on “#selfie” during each month based on every year of collected data.

The data extracted from “#selfie” suggests that there is an annual increase in the relative number of Pro-ED images posted during late spring and early summer (April to June). Furthermore, the data indicates that the relative number of Pro-ED images posted on “#selfie” reaches its highest value in the late summer or early autumn (July to September) every year, as is shown in Figure 3. Moreover, the data implies that the relative amount of Pro-ED posts on “#selfie” has been gradually increasing every year since 2017. The mean and standard deviation of the relative number of Pro-ED posts have also been growing since 2017. The growth in mean and standard deviation stipulates that over the last five years, the percentage, relative to the whole, of Pro-ED content has been increasing, and that the difference between the relative total of Pro-ED content in the summer and during other times of year has been growing as well. Based on the polynomial and the linear regressions, it is likely that this growth will continue into the next year and beyond.

## 5 Discussion

### 5.1 Principal Findings

This study explored the efficacy of training a deep learning model to classify Pro-ED content on the social media platform Twitter. Two computer vision models, ViT and ResNet, were trained using scraped tweets and then assessed to determine their accuracy on a hold-out test set. The results of this analysis indicated that both models were highly effective at recognizing Pro-ED content. Though both models do exceedingly well at recognizing Pro-ED content, they are both hindered by the scope of the training data which causes them to generate false positive classifications. One likely explanation for the higher number of false positive classifications in comparison to false negative classifications is

the presence of Pro-ED content within the Not Pro-ED set and test set. Such data overlap would cause a greater number of false positives and is indeed possible due to the unedited nature of the dataset.

Next, after training, the ViT model was deployed to analyze a vast number of images from Twitter’s “#selfie”. The results of the ViT’s analysis affirmed long-standing theories which relate worsening eating disorder symptoms to specific seasons of the year. Previous research has indicated that those who suffer from certain eating disorders—largely bulimia—see an annual increase in eating disorder severity during certain seasons (Fornari et al., 1994; Ghadirian et al., 1999; Yamatsuji et al., 2003). These seasonal changes in symptom intensity are often attributed to Seasonal Affective Disorder (SAD), which is a disorder characterized by seasonal patterns of major depression followed by a period of remission (Hardin et al., 1991; Nasteski, 2017). Though it was previously believed that SAD was reserved solely to winter, recent research has uncovered that for those suffering from body-image concerns, it may also occur during the summer (Akram et al., 2019; Ghadirian et al., 1999; Liang et al., 2018).

The analysis of “#selfie” illustrated similar seasonal fluctuations in Pro-ED sentiment—measured by the relative amount of Pro-ED content—to those described in studies that analyzed the seasonality of eating disorder symptoms (Ghadirian et al., 1999; Hardin et al., 1991; Melrose, 2015). In this instance, during late spring and early summer, the relative number of Pro-ED images on “#selfie” increases, reaching its peak during July or August. Afterward, the relative abundance of Pro-ED images falls, reaching its lowest point in autumn or early winter. This wave-like pattern consisting of an annual rise and fall in the relative number of Pro-ED content on “#selfie” is analogous to the seasonal fluctuations of SAD and is supported by research which shows that eating disorder symptoms, and, in turn, Pro-ED sentiment increases during the summer (Sloan, 2002). Moreover, a previous study determined that people’s mental well-being deteriorates during the summer and that social media reflects this deterioration through an increase in the amount of depressive language in posts (Burke et al., 2018). Since eating disorders and depression are closely linked, worsening symptoms of depression during the summer could also bring about an increase in eating disorder severity (Garcia et al., 2020; Mischoulon et al., 2011; Presnell et al., 2009; Sander et al., 2021), and therefore, dramatically enlarge the relative abundance of Pro-ED content, as is seen on “#selfie”.

Recent research has shown that the Covid-19 pandemic worsened the condition of eating disorder patients and led to an increase in new eating disorder cases (Khosravi, 2020; Rodgers et al., 2020; Taquet et al., 2022). The foregoing data from “#selfie” depicts a similar story. In January 2020, which marked the beginning of the global pandemic, there was a sudden increase in the relative amount of Pro-ED content. However, this sudden increase continued, and the number of Pro-ED images kept on rising all

through spring and summer. By August 2020, the relative number of Pro-ED posts on “#selfie” had outgrown all years prior. To this day, the year 2020 boasts the highest mean Pro-ED content of all time on “#selfie”. Despite the end of the global pandemic, the relative number of Pro-ED images on “#selfie” has not declined to its pre-pandemic levels. Moreover, the data does not indicate any future decline.

The ViT model's classifications of over four years' worth of images extracted from “#selfie” illustrated seasonal trends in content promoting eating disorders and an overall increase in said content over the last four years. These findings are analogous to the trends uncovered in previous studies on both the seasonality of eating disorders and the effect of social media on the spread of Pro-ED content. The congruency between the results of the ViT model and non-artificial intelligence-based studies of the same phenomenon indicates that the model can recognize and understand the nuances of Pro-ED images, differentiating it from other image types, even when analyzing swathes of data. Moreover, the ViT model reaches accuracies comparable to NLP-based models developed in previous studies, even though the ViT model analyzes an entirely different informative medium (Benítez-Andrades et al., 2021; Oksanen et al., 2015). Most importantly, like other artificial intelligence models, the model developed in this study can analyze images at speeds and accuracies that humans are unable to replicate, offering a quick and efficient method for recognizing Pro-ED content on social media.

There does not exist a perfect solution for halting the spread of Pro-ED content on social media. However, computer vision models, such as ViT, and NLP models, such as BERT, offer a possible countermeasure (Astorino et al., 2020; Benítez-Andrades et al., 2022). These models can be employed by social media companies to detect and remove content that the models classify as Pro-ED, or they can be used by external organizations to monitor the prevalence of Pro-ED content on social media. In any case, these models offer an efficient and accurate method of detecting Pro-ED content on social media and may be used to mitigate the spread of the Pro-ED movement, thereby restricting the harmful effects of such posts (Feldhege et al., 2021).

## 5.2 Limitations

There are several limitations in the research conducted. First, the dataset is comprised of images from a single social media platform. Second, the images in the datasets were labeled based on the hashtag from which they were downloaded and, therefore, there are potentially images that have an incorrect classification. Third, the Pro-ED hashtags that were scraped to build the dataset are comprised only of those Pro-ED hashtags that have been documented to support eating disorders. Images from many other Pro-ED hashtags have been omitted from the dataset.

## 6 Conclusion

Deep Learning computer vision models were trained to classify images extracted from Tweets as either promoting eating disorders or not. Both models achieved accuracies of greater than 83%. The best performing model was the Visual Transformer model, a model which employs a transformer-based architecture to generate a classification based on an input image.

The Visual Transformer model was used to analyze tweets that were posted on “#selfie” over the last five years. The analysis uncovered seasonal patterns of increased pro-eating disorder presence on social media during the summer season. Thus, providing much-needed empirical evidence to support earlier theories about seasonal fluctuations in eating disorder symptoms. Moreover, the Visual Transformer model also revealed the profound effect the Covid-19 pandemic has had on the prevalence of Pro-ED content on social media and revealed how the presence of stressors influences the abundance of Pro-ED content.

Future research will focus on increasing the current dataset to elevate model generalizability and erase any incidental bias, pairing image processing and natural language processing techniques to create a model capable of analyzing any social media post, and utilizing the trained model in the real world, potentially as a bot, to uncover and report posts and platforms that promote eating disorders.

## References

- A Hash Based Image Matching Algorithm for Social Networks* / SpringerLink. (n.d.). Retrieved September 2, 2022, from [https://link.springer.com/chapter/10.1007/978-3-319-61578-3\\_18](https://link.springer.com/chapter/10.1007/978-3-319-61578-3_18)
- Agostino, H., Burstein, B., Moubayed, D., Taddeo, D., Grady, R., Vyver, E., Dimitropoulos, G., Dominic, A., & Coelho, J. S. (2021). Trends in the Incidence of New-Onset Anorexia Nervosa and Atypical Anorexia Nervosa Among Youth During the COVID-19 Pandemic in Canada. *JAMA Network Open*, 4(12), e2137395. <https://doi.org/10.1001/jamanetworkopen.2021.37395>
- Akram, F., Jennings, T. B., Stiller, J. W., Lowry, C. A., & Postolache, T. T. (2019). Mood Worsening on Days with High Pollen Counts is associated with a Summer Pattern of Seasonality. *Pteridines*, 30(1), 133–141. <https://doi.org/10.1515/pteridines-2019-0016>
- Application of image hash algorithm in copyright protection system*. (n.d.). Retrieved September 2, 2022, from <https://www.spiedigitallibrary.org/conference-proceedings-of-spie/12167/121673B/Application-of-image-hash-algorithm-in-copyright-protection-system/10.1117/12.2628802.short?SSO=1&tab=ArticleLink>
- Arseniev-Koehler, A., Lee, H., McCormick, T., & Moreno, M. A. (2016). #Proana: Pro-Eating Disorder Socialization on Twitter. *Journal of Adolescent Health*, 58(6), 659–664. <https://doi.org/10.1016/j.jadohealth.2016.02.012>
- Astorino, A., Berti, R., Astorino, A., Bitonti, V., De Marco, M., Feraco, V., Palumbo, A., Porti, F., & Zannino, I. (2020). Early Detection of Eating Disorders Through Machine Learning Techniques. In I. S. Kotsireas & P. M. Pardalos (Eds.), *Learning and Intelligent Optimization* (pp. 33–39). Springer International Publishing. [https://doi.org/10.1007/978-3-030-53552-0\\_5](https://doi.org/10.1007/978-3-030-53552-0_5)
- Benítez-Andrades, J. A., Alija-Pérez, J.-M., Vidal, M.-E., Pastor-Vargas, R., & García-Ordás, M. T. (2022). Traditional Machine Learning Models and Bidirectional Encoder Representations From Transformer (BERT)–Based Automatic Classification of Tweets About Eating Disorders: Algorithm Development and Validation Study. *JMIR Medical Informatics*, 10(2), e34492. <https://doi.org/10.2196/34492>
- Bert, F., Gualano, M., Camussi, E., & Siliquini, R. (2016). Risks and Threats of Social Media Websites: Twitter and the Proana Movement. *Cyberpsychology, Behavior, and Social Networking*, 19. <https://doi.org/10.1089/cyber.2015.0553>
- BERT Model-Based Approach For Detecting Categories of Tweets in the Field of Eating Disorders (ED)* / IEEE Conference Publication / IEEE Xplore. (n.d.). Retrieved September 2, 2022, from <https://ieeexplore.ieee.org/document/9474757>
- Branley, D. B., & Covey, J. (2017). Pro-ana versus Pro-recovery: A Content Analytic Comparison of Social Media Users' Communication about

- Eating Disorders on Twitter and Tumblr. *Frontiers in Psychology*, 8, 1356. <https://doi.org/10.3389/fpsyg.2017.01356>
- Burke, M., González, F., Baylis, P., Heft-Neal, S., Baysan, C., Basu, S., & Hsiang, S. (2018). Higher temperatures increase suicide rates in the United States and Mexico. *Nature Climate Change*, 8(8), Article 8. <https://doi.org/10.1038/s41558-018-0222-x>
- Chen, X., Hsieh, C.-J., & Gong, B. (2022). *When Vision Transformers Outperform ResNets without Pre-training or Strong Data Augmentations* (arXiv:2106.01548). arXiv. <http://arxiv.org/abs/2106.01548>
- Culbert, K. M., Racine, S. E., & Klump, K. L. (2015). Research Review: What we have learned about the causes of eating disorders – a synthesis of sociocultural, psychological, and biological research. *Journal of Child Psychology and Psychiatry*, 56(11), 1141–1164. <https://doi.org/10.1111/jcpp.12441>
- Deng, J., Dong, W., Socher, R., Li, L.-J., Li, K., & Fei-Fei, L. (n.d.). *ImageNet: A Large-Scale Hierarchical Image Database*. 8.
- Dignard, N. A. L., & Jarry, J. L. (2021). The “Little Red Riding Hood effect:” Fitspiration is just as bad as thinspiration for women’s body satisfaction. *Body Image*, 36, 201–213. <https://doi.org/10.1016/j.bodyim.2020.11.012>
- Dosovitskiy, A., Beyer, L., Kolesnikov, A., Weissenborn, D., Zhai, X., Unterthiner, T., Dehghani, M., Minderer, M., Heigold, G., Gelly, S., Uszkoreit, J., & Houlsby, N. (2021). *An Image is Worth 16x16 Words: Transformers for Image Recognition at Scale* (arXiv:2010.11929). arXiv. <http://arxiv.org/abs/2010.11929>
- Fardouly, J., Crosby, R. D., & Sukunesan, S. (2022). Potential benefits and limitations of machine learning in the field of eating disorders: Current research and future directions. *Journal of Eating Disorders*, 10(1), 66. <https://doi.org/10.1186/s40337-022-00581-2>
- Feldhege, J., Moessner, M., & Bauer, S. (2021). Detrimental Effects of Online Pro-Eating Disorder Communities on Weight Loss and Desired Weight: Longitudinal Observational Study. *Journal of Medical Internet Research*, 23(10), e27153. <https://doi.org/10.2196/27153>
- Fornari, V. M., Braun, D. L., Sunday, S. R., Sandberg, D. E., Matthews, M., Chen, li-L., Mandel, F. S., Halmi, K. A., & Katz, J. L. (1994). Seasonal patterns in eating disorder subgroups. *Comprehensive Psychiatry*, 35(6), 450–456. [https://doi.org/10.1016/0010-440X\(94\)90228-3](https://doi.org/10.1016/0010-440X(94)90228-3)
- Garcia, S. C., Mikhail, M. E., Keel, P. K., Burt, S. A., Neale, M. C., Boker, S., & Klump, K. L. (2020). Increased rates of eating disorders and their symptoms in women with major depressive disorder and anxiety disorders. *The International Journal of Eating Disorders*, 53(11), 1844–1854. <https://doi.org/10.1002/eat.23366>
- Ghadirian, A. M., Marini, N., Jabalpurwala, S., & Steiger, H. (1999). Seasonal mood patterns in eating disorders. *General Hospital Psychiatry*, 21(5), 354–359. [https://doi.org/10.1016/s0163-8343\(99\)00028-6](https://doi.org/10.1016/s0163-8343(99)00028-6)
- Goh, Y. C., Cai, X. Q., Theseira, W., Ko, G., & Khor, K. A. (2020). Evaluating human versus machine learning performance in classifying

- research abstracts. *Scientometrics*, 125(2), 1197–1212.  
<https://doi.org/10.1007/s11192-020-03614-2>
- Gordon, R. A. (2000). *Eating disorders: Anatomy of a social epidemic, 2nd ed* (pp. xvi, 281). Blackwell Publishing.
- Han, S., Park, G., Lim, W., Kim, M., Na, J.-I., Park, I., & Chang, S. (2018). Deep neural networks show an equivalent and often superior performance to dermatologists in onychomycosis diagnosis: Automatic construction of onychomycosis datasets by region-based convolutional deep neural network. *PLOS ONE*, 13, e0191493.  
<https://doi.org/10.1371/journal.pone.0191493>
- Hardin, T. A., Wehr, T. A., Brewerton, T., Kasper, S., Berrettini, W., Rabkin, J., & Rosenthal, N. E. (1991). Evaluation of seasonality in six clinical populations and two normal populations. *Journal of Psychiatric Research*, 25(3), 75–87. [https://doi.org/10.1016/0022-3956\(91\)90001-Q](https://doi.org/10.1016/0022-3956(91)90001-Q)
- He, K., Zhang, X., Ren, S., & Sun, J. (2015). *Deep Residual Learning for Image Recognition* (arXiv:1512.03385). arXiv.  
<http://arxiv.org/abs/1512.03385>
- Howell, C. R., Su, W., Nassel, A. F., Agne, A. A., & Cherrington, A. L. (2020). Area based stratified random sampling using geospatial technology in a community-based survey. *BMC Public Health*, 20(1), 1678. <https://doi.org/10.1186/s12889-020-09793-0>
- Islam, K. (2022). *Recent Advances in Vision Transformer: A Survey and Outlook of Recent Work* (arXiv:2203.01536). arXiv.  
<http://arxiv.org/abs/2203.01536>
- Jeeveswaran, K., Kathiresan, S., Varma, A., Magdy, O., Zonooz, B., & Arani, E. (2022). *A Comprehensive Study of Vision Transformers on Dense Prediction Tasks* (arXiv:2201.08683). arXiv.  
<http://arxiv.org/abs/2201.08683>
- Jiang, T., Gradus, J. L., & Rosellini, A. J. (2020). Supervised machine learning: A brief primer. *Behavior Therapy*, 51(5), 675–687.  
<https://doi.org/10.1016/j.beth.2020.05.002>
- Khosravi, M. (2020). The challenges ahead for patients with feeding and eating disorders during the COVID-19 pandemic. *Journal of Eating Disorders*, 8(1), 43. <https://doi.org/10.1186/s40337-020-00322-3>
- Kim, J., Lee, J., Park, E., & Han, J. (2020). A deep learning model for detecting mental illness from user content on social media. *Scientific Reports*, 10(1), Article 1. <https://doi.org/10.1038/s41598-020-68764-y>
- Li, S., Jiao, J., Han, Y., & Weissman, T. (2017). *Demystifying ResNet* (arXiv:1611.01186). arXiv. <https://doi.org/10.48550/arXiv.1611.01186>
- Liang, C.-S., Chung, C.-H., Tsai, C.-K., & Chien, W.-C. (2018). Seasonality of hospital admissions and birth dates among inpatients with eating disorders: A nationwide population-based retrospective study. *Eating and Weight Disorders: EWD*, 23(2), 233–240. <https://doi.org/10.1007/s40519-016-0326-0>



- M, T., O, C., L, M., & Z, B. (2018). Tweeting weight loss: A comparison of #thinspiration and #fitspiration communities on Twitter. *Body Image*, 25. <https://doi.org/10.1016/j.bodyim.2018.03.002>
- Matthews, A., Kramer, R. A., Peterson, C. M., & Mitan, L. (2021). Higher admission and rapid readmission rates among medically hospitalized youth with anorexia nervosa/atypical anorexia nervosa during COVID-19. *Eating Behaviors*, 43, 101573. <https://doi.org/10.1016/j.eatbeh.2021.101573>
- Melrose, S. (2015). Seasonal Affective Disorder: An Overview of Assessment and Treatment Approaches. *Depression Research and Treatment*, 2015, 178564. <https://doi.org/10.1155/2015/178564>
- Mento, C., Silvestri, M. C., Muscatello, M. R. A., Rizzo, A., Celebre, L., Praticò, M., Zoccali, R. A., & Bruno, A. (2021). Psychological Impact of Pro-Anorexia and Pro-Eating Disorder Websites on Adolescent Females: A Systematic Review. *International Journal of Environmental Research and Public Health*, 18(4), 2186. <https://doi.org/10.3390/ijerph18042186>
- Mischoulon, D., Eddy, K. T., Keshaviah, A., Dinescu, D., Ross, S. L., Kass, A. E., Franko, D. L., & Herzog, D. B. (2011). Depression and Eating Disorders: Treatment and Course. *Journal of Affective Disorders*, 130(3), 470–477. <https://doi.org/10.1016/j.jad.2010.10.043>
- Nasteski, V. (2017). An overview of the supervised machine learning methods. *HORIZONS.B*, 4, 51–62. <https://doi.org/10.20544/HORIZONS.B.04.1.17.P05>
- Oksanen, A., Garcia, D., Sirola, A., Näsi, M., Kaakinen, M., Keipi, T., & Räsänen, P. (2015). Pro-Anorexia and Anti-Pro-Anorexia Videos on YouTube: Sentiment Analysis of User Responses. *Journal of Medical Internet Research*, 17(11), e5007. <https://doi.org/10.2196/jmir.5007>
- Otto, A. K., Jary, J. M., Sturza, J., Miller, C. A., Prohaska, N., Bravender, T., & Van Huisse, J. (2021). Medical Admissions Among Adolescents With Eating Disorders During the COVID-19 Pandemic. *Pediatrics*, 148(4), e2021052201. <https://doi.org/10.1542/peds.2021-052201>
- Parsons, H., Murphy, B., Malone, D., & Holme, I. (2021). Review of Ireland’s First Year of the COVID-19 Pandemic Impact on People Affected by Eating Disorders: ‘Behind Every Screen There Was a Family Supporting a Person with an Eating Disorder.’ *Journal of Clinical Medicine*, 10(15), Article 15. <https://doi.org/10.3390/jcm10153385>
- Presnell, K., Stice, E., Seidel, A., & Madeley, M. C. (2009). Depression and Eating Pathology: Prospective Reciprocal Relations in Adolescents. *Clinical Psychology & Psychotherapy*, 16(4), 357–365. <https://doi.org/10.1002/cpp.630>
- Rao, A., Park, J., Woo, S., Lee, J.-Y., & Aalami, O. (2021). *Studying the Effects of Self-Attention for Medical Image Analysis* (arXiv:2109.01486). arXiv. <http://arxiv.org/abs/2109.01486>
- Ren, Y., Lu, C., Yang, H., Ma, Q., Barnhart, W. R., Zhou, J., & He, J. (2022). Using machine learning to explore core risk factors associated with the risk of eating disorders among non-clinical young women in China: A

- decision-tree classification analysis. *Journal of Eating Disorders*, 10(1), 19. <https://doi.org/10.1186/s40337-022-00545-6>
- Rodgers, R. F., Lombardo, C., Cerolini, S., Franko, D. L., Omori, M., Fuller-Tyszkiewicz, M., Linardon, J., Courtet, P., & Guillaume, S. (2020). The impact of the COVID-19 pandemic on eating disorder risk and symptoms. *International Journal of Eating Disorders*, 53(7), 1166–1170. <https://doi.org/10.1002/eat.23318>
- Rosen, A. O., Holmes, A. L., Balluerka, N., Hidalgo, M. D., Gorostiaga, A., Gómez-Benito, J., & Huedo-Medina, T. B. (2022). Is Social Media a New Type of Social Support? Social Media Use in Spain during the COVID-19 Pandemic: A Mixed Methods Study. *International Journal of Environmental Research and Public Health*, 19(7), Article 7. <https://doi.org/10.3390/ijerph19073952>
- Sander, J., Moessner, M., & Bauer, S. (2021). Depression, Anxiety and Eating Disorder-Related Impairment: Moderators in Female Adolescents and Young Adults. *International Journal of Environmental Research and Public Health*, 18(5), 2779. <https://doi.org/10.3390/ijerph18052779>
- Shadish, W. R., Cook, T. D., & Campbell, D. T. (2002). *Experimental and quasi-experimental designs for generalized causal inference* (pp. xxi, 623). Houghton, Mifflin and Company.
- Sloan, D. M. (2002). Does warm weather climate affect eating disorder pathology? *International Journal of Eating Disorders*, 32(2), 240–244. <https://doi.org/10.1002/eat.10077>
- Spettigue, W., & Henderson, K. A. (2004). Eating Disorders and the Role of the Media. *The Canadian Child and Adolescent Psychiatry Review*, 13(1), 16–19.
- Sukunesan, S., Huynh, M., & Sharp, G. (2021). Examining the Pro-Eating Disorders Community on Twitter Via the Hashtag #proana: Statistical Modeling Approach. *JMIR Mental Health*, 8(7), e24340. <https://doi.org/10.2196/24340>
- Tan, C., Sun, F., Kong, T., Zhang, W., Yang, C., & Liu, C. (2018). *A Survey on Deep Transfer Learning* (arXiv:1808.01974). arXiv. <https://doi.org/10.48550/arXiv.1808.01974>
- Taquet, M., Geddes, J. R., Luciano, S., & Harrison, P. J. (2022). Incidence and outcomes of eating disorders during the COVID-19 pandemic. *The British Journal of Psychiatry*, 220(5), 262–264. <https://doi.org/10.1192/bjp.2021.105>
- Targ, S., Almeida, D., & Lyman, K. (2016). *Resnet in Resnet: Generalizing Residual Architectures* (arXiv:1603.08029). arXiv. <https://doi.org/10.48550/arXiv.1603.08029>
- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, L., & Polosukhin, I. (2017). *Attention Is All You Need* (arXiv:1706.03762). arXiv. <https://doi.org/10.48550/arXiv.1706.03762>
- Wu, Z., Shen, C., & van den Hengel, A. (2019). Wider or Deeper: Revisiting the ResNet Model for Visual Recognition. *Pattern Recognition*, 90, 119–133. <https://doi.org/10.1016/j.patcog.2019.01.006>

- Yamatsuji, M., Yamashita, T., Arie, I., Taga, C., Tatara, N., & Fukui, K. (2003). Seasonal variations in eating disorder subtypes in Japan. *International Journal of Eating Disorders*, 33(1), 71–77. <https://doi.org/10.1002/eat.10107>
- Yu, J., Wang, Z., Vasudevan, V., Yeung, L., Seyedhosseini, M., & Wu, Y. (2022). *CoCa: Contrastive Captioners are Image-Text Foundation Models* (arXiv:2205.01917). arXiv. <https://doi.org/10.48550/arXiv.2205.01917>
- Zhou, Z., Shin, J., Zhang, L., Gurudu, S., Gotway, M., & Liang, J. (2017). Fine-tuning Convolutional Neural Networks for Biomedical Image Analysis: Actively and Incrementally. *Proceedings. IEEE Computer Society Conference on Computer Vision and Pattern Recognition, 2017*, 4761–4772. <https://doi.org/10.1109/CVPR.2017.506>