

Segmented Image Compression in Healthcare

Alex Nava, Cristina Bonilla Bernal, Jayden Tang, Logan Graves

Mentors: Ayushman Chakraborty, Qingxi Meng

Abstract

The crossroads at which medical imaging and data compression intersect has yielded a fascinating, novel area of research, particularly pertaining to the Segment Anything Model (SAM), an AI-based image segmentation model. We researched a plethora of standard medical imaging techniques, including Computed Tomography (CT scans), Positron Emission Tomography (PET scans), Ultrasound, and Magnetic Resonance Imaging (MRI). Additionally, we analyzed more specific areas of medical imaging such as digital pathology, mammography, and photoacoustic imaging. To supplement our knowledge of the different types of imaging, we researched, specifically, how MRI scans are processed in terms of segmentation and standard storage, and compression practices. Furthermore, we studied recurrent difficulties that medical professionals face when segmenting certain areas of the body, paying particular attention to issues within the brain and the spinal cord. By speaking to a member of the Radiology Interest Group at Stanford, we also determined frequent issues surrounding storage and clinical workflow, thus narrowing our research into how the Segment Anything Model can be applied in a robust, efficient, and critical manner.

Integrating our knowledge about various kinds of medical imaging technology, we present a proof concept for a novel image compression technique based on SAM, one which is especially suited to medical imaging technology. By automatically distinguishing between unimportant image aspects (such as the blank black background of an MRI) and important aspects (such as the anatomical details of the scan), we can apply lossy compression to nonessential aspects and lossless compression to essential ones, allowing much greater amounts of compression without losing details relevant to the scans. We compare this technique with existing compression methods and suggest its potential applications, as well as areas for future research.

Background

Medical imaging maintains an invaluable role in the healthcare field, advancing the processes of diagnosis, treatment, and recovery in ways that go beyond human ability. A plethora of imaging techniques exist; however, the most prevalent prove to be X-ray imaging, Computed Tomography (CT) imaging, and Magnetic Resonance Imaging (MRI). All of these techniques, despite their inherent differences, are linked by a common thread: the necessity for digital storage. Digital storage is defined by the process of storing and retaining information through binary code, allowing for the preservation of images, videos, text, and other forms of digital data. Within the medical field, the transition from film-based imaging to digital storage has proven to be revolutionary; enhancements to image quality, accessibility, long-term retention, and interdisciplinary collaboration are significant. The coupling of medical imaging and the digital world requires a closer look into just how much storage these images utilize, particularly into the role that data compression plays. Data compression is the process by which digital information is encoded to use fewer bits to represent the same information. The compression of digital data optimizes storage space, transmission times, and transfer efficiency, reflecting its gravity within the healthcare field.

Historical Context

The transition from film-based imaging to digital imaging in the 1980s and 1990s marked the beginning of a large area of focus on data compression in the medical field. Early compression techniques, mainly lossless compression algorithms, were employed to reduce the amount of storage taken up by patient records and images. In the late 1990s and early 2000s, telemedicine saw an increase in popularity, thus emphasizing the necessity of data compression as doctors required access to medical images. The advent of cloud storage and its worldwide proliferation in the early 2000s furthered the necessity for efficient data storage, underscoring the value of data compression research in healthcare [1]. Soon after, DICOM (Digital Imaging and Communications in Medicine) became the international standard, defining specific guidelines to maximize efficiency and collaboration when analyzing medical information. These guidelines popularized JPEG and JPEG2000 compression formats, allowing for easy, efficient access and storage to medical images [2]. In recent years, however, research into machine learning and its role in data compression has become increasingly popular, placing a strong emphasis on artificial intelligence (AI) and the development of compression algorithms. Considering this historical context, our present study explored a novel compression technique that utilizes the Segment Anything Model (SAM), and by

building off historical milestones in the field of data compression, we developed a proof of concept compression technique.

Segment Anything Model (SAM)

The Segment Anything Model (SAM) is a segmentation system that runs entirely off user prompts, relying on its zero-shot generalization feature to identify and classify unknown objects and images. Without any need for additional training, SAM is able to process a vast array of user input prompts, and when given a grid of points, is able to segment out certain areas of an image. SAM's dataset consists of over 1.1 billion segmentation masks which were derived from ~11 million images, demonstrating its robust accuracy and efficiency. Additionally, SAM was purposefully decoupled into a one time image-encoder and a lightweight mask decoder, enabling the model to simply be run on a web browser [3]. Overall, SAM's proficiency in segmenting medical images, specifically MRI scans, played a vital role in our research.

Theory and Research Objectives

Our team theorized that by compressing segments made by SAM individually, we would achieve higher rates of compression than if we compressed the image as a whole, as is done traditionally. This essentially means that by compressing the disparate parts of an MRI scan, the background and foreground, separately, results would yield a greater rate of compression. To test this theory, we split our objectives into two areas of focus: medical imaging/SAM and the development of our model. Our team researched the intricacies and applications of medical imaging, specifically Magnetic resonance imaging (MRI). Our goal was to build a wealth of knowledge that would enable us to develop a model that could be effectively employed in the medical field; in other words, we intended to become experts in the field of medical imaging so as to apply our knowledge accurately when developing our compression algorithm. Our ultimate objective, backed by the extensive research into medical imaging and SAM, was to create a proof-of-concept image compression technique, one that utilized our original theory of individual segmentation and compression.

Significance and Application

The significance of our research goes beyond the rates of compression we achieved, rather, it is most effectively applied when handling storage costs of medical images. Maintaining an information system capable of storing medical images for a prolonged period of time, a minimum of seven years in California

for example, comes at a high price [4]. Storage costs generally range from \$25,000,000 per year, proving to be a major hindrance to accessibility and affordability [5]. For well-funded hospitals located in high-income areas, digital storage costs aren't that large of an issue; however, for hospitals that are found in low-income areas without major funding, these costs are a major obstacle. Hospitals that fit this description are forced onto a tight budget, meaning that their services and functioning as a whole lack behind their well-funded counterparts. By achieving rates of compression greater than 50% on MRI scans, this effectively means that storage costs will be reduced extensively. This reduction in cost allows for an increased number of hospitals and radiology practices to be established in low-income areas as the overall cost of running the practice decreases.

Methods

Medical imaging provides a uniquely suitable forum and use case for building and testing prototypes: not only are image storage costs a significant burden on many private practices, and cause a significant amount of carbon emissions, but these images are also much simpler in content than most other images, such as those taken with regular cameras. Most medical imaging devices output relatively simple images: in the case of MRIs, this means a black-and-white image of the targeted anatomical structures, set on a simple black background. Because the background and foreground layers are so distinct, this makes them relatively trivial to segment, eliminating the need to do particular fine-tuning or prompting on SAM and allowing us to focus on the algorithmic compression aspects of the project rather than the particulars of the image segmentation technology.

We designed and coded an algorithm in Python, utilizing a number of well-known libraries such as Pytorch for our model implementation, Numpy and matplotlib for data visualization, and PIL (python image library) for certain image display and conversion functions. In addition to these python packages, we wrote adaptable commands for Imagemagick, a command-line utility for image manipulation and conversion.

This algorithm consisted of the following steps:

1. Set up SAM(download and install the open-weights model)
2. Iterate through a given directory of images, or a single image, doing the following:

- a. Convert the files to a common format (usually lossless JPEG or PNG; can be customized to use case, but doesn't matter much in final results)
- b. Use SAM to create image masks separating foreground from background
- c. Select the mask that correctly segments the foreground from background
- d. Split the background off from the main image, and apply extremely lossy JPEG compression to it (quality 1, i.e. maximum compression available to the format)
- e. Keep foreground in a lossless format (PNG)

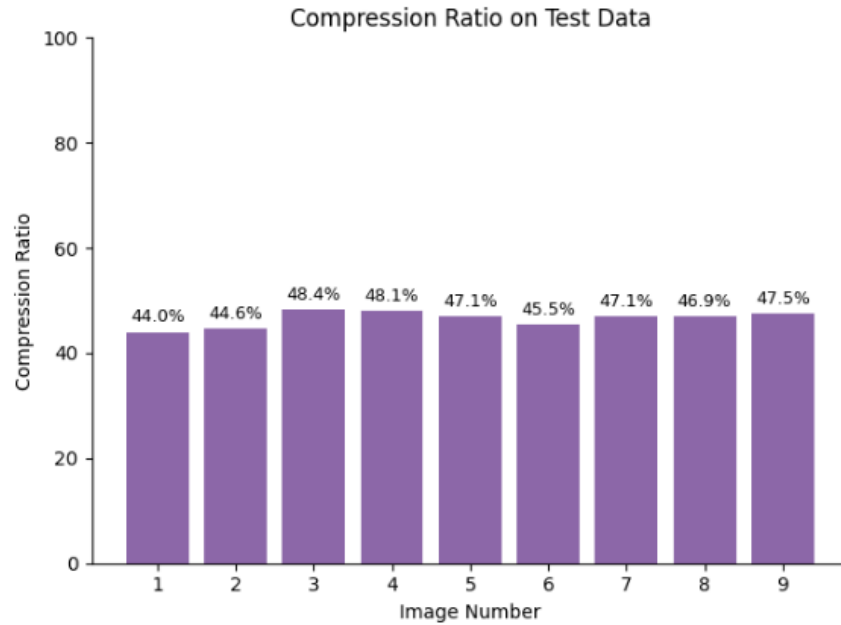
Images could at this point be stored separately, which could in theory lead to more significant compression results if implemented correctly. However, we chose to instead recombine the images into a lossless format.

- f. Recombine images to obtain an image with resolution equal to the resolution of the input image

3. Output final image

The algorithm also can optionally output intermediate images, such as the background and foreground layers. Its source code is available as a Jupyter Notebook, downloadable from the Google colab available at <https://colab.research.google.com/drive/lzfXycr4ULxocHKvjPtvnHDsGi3MUa-w>.

This algorithm is comfortable to use for those familiar with its basic structure and those familiar with the command-line interface. However, we expect that many doctors and other medical professionals may *not* be either one of those two, and therefore find it cumbersome or even impossible to use in its basic code form. Using Streamlit, a novel python framework, we created a web app to enable use of the algorithm simply and visually. Our web app prompts the user to upload a zipped dataset of images which will be altered with our algorithm. With the given dataset, the user will be shown a number range slider, which will then allow them to compress an image of their choice. The user will be shown the corresponding original image and altered image side by side relative to the index of the slider. The accessibility of our web app makes it easy to navigate through a whole dataset of images and use our algorithm with said dataset.



A visualization of compression results on one test datasets, a set of layers of a skull MRI.




Image demonstrating our web app with use of the number slider and algorithm on a image of a brain MRI

Results

The given outcomes of our proof-of-concept generalized algorithm that utilizes image segmentation technology for compression results yielded positive compression results. Our intelligent compression algorithm achieved greater than 50% compression on MRI tests and greater than 10% compression with other images. Comparing original MRI images and altered images that used our algorithms results in

visually lossless results; in the general use case, results are not visually lossless (because medical images are much more complex and the background is not near-uniform or uniform to begin with). Though the inner workings of compression algorithms can be complex and unclear, we believe that these results are possible because the applied JPEG compression makes the background layers extremely easy to store in comparison to the original image by making them almost uniform pixel values; because common compression algorithms use similar techniques (such as dictionary encoding with Huffman coding) the resulting image's space savings are transferable once the lossy compression has been applied and converted into a lossless image. .

Conclusions

Based on our positive results, we can conclude that SAM's application in the field of data compression and medical imaging has promising results for future studies. Applications of our model in the medical field are vast and auspicious, displaying promise in areas of data transmission, storage, and various other methods. Additionally, our model is essentially applicable anywhere where regular lossy compression is, and in some cases where lossless compression is. With more precise fine-tuning, or different image segmentation technology, we believe that much greater than 10% savings could be achieved on the general use case, though without visually lossless results. Furthermore, our model can be expanded upon by means of a classification system. Once segmented, the model could be taught to label anatomical structures, identifying respective parts of a medical image. Although, this would require an extensive research, it is a viable and invaluable asset to the medical field. Ultimately, our compression model has proven its success in reducing storage size in a way that maintains the original quality and data, making our proof of concept a practical compression technique that cuts costs and sets up imaging practices to be more affordable and cost effective.

Future Directions

One promising avenue for our future research involves an expansion of the compression framework to encompass general images. Although we anticipate that the compression ratios might not be as favorable and the model's performance may exhibit some degree of variability, this direction could yield broader applications. Diverse domains could benefit from the extended versatility of such an approach.

Furthermore, an intriguing trajectory to explore centers around adapting our work for efficient data transmission. Specifically, leveraging our compressed data transmission methodology could substantially enhance bandwidth efficiency. This, in turn, holds the potential to significantly reduce the time required for transferring medical images between endpoints. While our existing work concentrated on the distinction between background and foreground, an exciting possibility lies in delving into more intricate segmentation techniques. By achieving finer-grained segmentation, we could potentially unlock even greater gains in compression efficiency.

Beyond the realm of storage optimization, another compelling avenue emerges: the application of compression techniques for purposes beyond their conventional scope. For instance, an intriguing prospect lies in harnessing compression as a tool for quantifying discrepancies within medical imaging. This innovative direction holds the promise of not only enhancing our understanding of these discrepancies but also advancing their diagnostic potential for future studies that can increase the potential of developing a greater and better access to medical treatments.

Furthermore, by enhancing the potential to transfer and store medical imaging data, it can be substantially used to bring it to communities where access to different medical studies is reduced due to the lack of information and specific tools to analyze medical imaging data information. Thus, by making it easier to transfer and store those images, there is a greater possibility of bringing the knowledge to places where medical information is not as easily accessible; potentially changing the perspective of different medical treatments used in those places and increasing their effectiveness.

In summary, our future directions span a spectrum of possibilities. From broadening the compression framework to accommodate general images and enhancing data transmission efficiency, to helping other communities with their medical knowledge and access, to pursuing intricate segmentation methodologies and repurposing compression for novel analytical insights. Each avenue holds the potential to contribute significantly to both the field of image compression and medical imaging as a whole. As we embark on these exciting trajectories, we look forward to pushing the boundaries of what is achievable and making meaningful strides in these dynamic areas of research.

References

[1]<https://www.ncbi.nlm.nih.gov/books/NBK207141/>

[2]<https://pubmed.ncbi.nlm.nih.gov/8057942/>

[3]<https://segment-anything.com/>

[4]<https://www.mbc.ca.gov/FAQs/?cat=Consumer&topic=Complaint:%20Medical%20Records#:~:text=HSC%20section%20123145%20indicates%20that,following%20discharge%20of%20the%20patient.>

[5]<https://www.massdevice.com/medical-data-storage-adding-cost-digitizing-health-records/>