


New AI method captures uncertainty in medical images

April 11 2024, by Adam Zewe



Tyche: the first in-context stochastic segmentation framework. Human annotators (top) can handle a wide variety of tasks, and different annotators

often produce differing segmentations. Existing automated methods (middle) are typically task-specific and provide only one segmentation per image. Tyche (bottom) can capture the disagreement among annotators across many modalities and anatomies without retraining or fine-tuning. Credit: *arXiv* (2024). DOI: 10.48550/arxiv.2401.13650

In biomedicine, segmentation involves annotating pixels from an important structure in a medical image, like an organ or cell. Artificial intelligence models can help clinicians by highlighting pixels that may show signs of a certain disease or anomaly.

However, these models typically only provide one answer, while the problem of medical image [segmentation](#) is often far from black and white. Five expert human annotators might provide five different segmentations, perhaps disagreeing on the existence or extent of the borders of a nodule in a lung CT image.

"Having options can help in decision-making. Even just seeing that there is uncertainty in a medical image can influence someone's decisions, so it is important to take this uncertainty into account," says Marianne Rakic, an MIT computer science Ph.D. candidate.

Rakic is the lead author of a paper with others at MIT, the Broad Institute of MIT and Harvard, and Massachusetts General Hospital that introduces a new AI tool that can capture the uncertainty in a medical image.

Known as Tyche (named for the Greek divinity of chance), the system provides multiple plausible segmentations that each highlight slightly different areas of a medical image. A user can specify how many options Tyche outputs and select the most appropriate one for their purpose.

Importantly, Tyche can tackle new segmentation tasks without needing to be retrained. Training is a data-intensive process that involves showing a model many examples and requires extensive machine-learning experience.

Because it doesn't need retraining, Tyche could be easier for clinicians and [biomedical researchers](#) to use than some other methods. It could be applied "out of the box" for a variety of tasks, from identifying lesions in a lung X-ray to pinpointing anomalies in a brain MRI.

Ultimately, this system could improve diagnoses or aid in biomedical research by calling attention to potentially crucial information that other AI tools might miss.

"Ambiguity has been understudied. If your model completely misses a nodule that three experts say is there and two experts say is not, that is probably something you should pay attention to," adds senior author Adrian Dalca, an assistant professor at Harvard Medical School and MGH and a research scientist in the MIT Computer Science and Artificial Intelligence Laboratory (CSAIL).

Their co-authors include Hallee Wong, a graduate student in [electrical engineering](#) and computer science; Jose Javier Gonzalez Ortiz, Ph.D. '23; Beth Cimini, associate director for bioimage analysis at the Broad Institute; and John Guttag, the Dugald C. Jackson Professor of Computer Science and Electrical Engineering. Rakic will present Tyche at the IEEE Conference on Computer Vision and Pattern Recognition, where Tyche has been selected as a highlight.

Addressing ambiguity

AI systems for medical image segmentation typically use neural networks. Loosely based on the human brain, [neural networks](#) are

machine-learning models comprising many interconnected layers of nodes, or neurons, that process data.

After speaking with collaborators at the Broad Institute and MGH who use these systems, the researchers realized two major issues limit their effectiveness. The models cannot capture uncertainty, and they must be retrained for even a slightly different segmentation task.

Some methods try to overcome one pitfall, but tackling both problems with a single solution has proven especially tricky, Rakic says.

"If you want to take ambiguity into account, you often have to use an extremely complicated model. With the method we propose, our goal is to make it easy to use with a relatively small model so that it can make predictions quickly," she says.

The researchers built Tyche by modifying a straightforward neural network architecture.

A user first feeds Tyche a few examples that show the segmentation task. For instance, examples could include several images of lesions in a heart MRI that different human experts have segmented so the model can learn the task and see that there is ambiguity.

The researchers found that just 16 example images, called a "context set," is enough for the model to make good predictions, but there is no limit to the number of examples one can use. The context set enables Tyche to solve new tasks without retraining.

For Tyche to capture uncertainty, the researchers modified the neural network so it outputs multiple predictions based on one medical image input and the context set. They adjusted the network's layers so that, as data moves from layer to layer, the candidate segmentations produced at

each step can "talk" to each other and the examples in the context set.

In this way, the model can ensure that candidate segmentations are all a bit different but still solve the task.

"It is like rolling dice. If your model can roll a two, three, or four but doesn't know you have a two and a four already, then either one might appear again," she says.

They also modified the training process so it is rewarded by maximizing the quality of its best prediction.

If the user asked for five predictions, at the end they can see all five medical image segmentations Tyche produced, even though one might be better than the others.

The researchers also developed a version of Tyche that can be used with an existing, pretrained model for medical image segmentation. In this case, Tyche enables the model to output multiple candidates by making slight transformations to images.

Better, faster predictions

When the researchers tested Tyche with datasets of annotated medical images, they found that its predictions captured the diversity of human annotators and that its best predictions were better than any from the baseline models. Tyche also performed faster than most models.

"Outputting multiple candidates and ensuring they are different from one another really gives you an edge," Rakic says.

The researchers also saw that Tyche could outperform more complex models that have been trained using a large, specialized dataset.

For future work, they plan to try using a more flexible context set, perhaps including text or multiple types of images. In addition, they want to explore methods that could improve Tyche's worst predictions and enhance the system so it can recommend the best segmentation candidates.

The research is [published](#) on the *arXiv* preprint server.

More information: Marianne Rakic et al, Tyche: Stochastic In-Context Learning for Medical Image Segmentation, *arXiv* (2024). [DOI: 10.48550/arxiv.2401.13650](https://doi.org/10.48550/arxiv.2401.13650)

This story is republished courtesy of MIT News (web.mit.edu/newsoffice/), a popular site that covers news about MIT research, innovation and teaching.

Provided by Massachusetts Institute of Technology

Citation: New AI method captures uncertainty in medical images (2024, April 11) retrieved 2 May 2024 from <https://medicalxpress.com/news/2024-04-ai-method-captures-uncertainty-medical.html>

<p>This document is subject to copyright. Apart from any fair dealing for the purpose of private study or research, no part may be reproduced without the written permission. The content is provided for information purposes only.</p>
--