

What is claimed is:

1. A computer-implemented method for processing electronic slide images corresponding to a tissue specimen, the method comprising:
 - receiving one or more electronic slide images associated with a tissue specimen, the tissue specimen being associated with a patient and/or medical case;
 - partitioning a first slide image of the one or more electronic slide images into a plurality of tiles;
 - detecting a plurality of tissue regions of the first slide image and/or plurality of tiles to generate a tissue mask;
 - determining whether any of the plurality of tiles corresponds to non-tissue;
 - removing any of the plurality of tiles that are determined to be non-tissue;
 - determining a prediction, using a machine learning prediction model, for at least one label for the one or more electronic slide images, the machine learning prediction model having been generated by processing a plurality of training images;
 - and
 - outputting the prediction of the trained machine learning prediction model.
2. The computer-implemented method of claim 1, wherein the plurality of tiles that are determined to be non-tissue are further determined to be a background of the tissue specimen.
3. The computer-implemented method of claim 1, wherein detecting the plurality of tissue regions comprises segmenting the tissue regions from a background of the one or more electronic slide images.

4. The computer-implemented method of claim 3, further comprising:
upon segmenting the tissue regions from the background, generating a tissue mask, the segmenting using thresholding based on color/intensity and/or texture features.

5. The computer-implemented method of claim 1, wherein the plurality of training images comprise a plurality of electronic slide images and a plurality of target labels.

6. The computer-implemented method of claim 1, wherein processing the plurality of training images comprises:

receiving a collection of digital images associated with at least one training tissue specimen, wherein the collection of digital images comprises a plurality of training electronic slide images;

receiving a plurality of synoptic annotations comprising one or more labels for each of the plurality of training electronic slide images;

partitioning one of the plurality of training electronic slide images into a plurality of training tiles for the plurality of training electronic slide images;

segmenting at least one tissue region from a background of the one or more electronic slide images to create a training tissue mask;

removing at least one of the plurality of tiles that are detected to be non-tissue; and

training the machine learning prediction model under weak supervision to infer at least one multi-label tile-level prediction using at least one label of the plurality of synoptic annotations.

7. The computer-implemented method of claim 6, wherein training the machine learning prediction model under weak supervision comprises using at least one of multiple-instance learning (MIL), Multiple Instance Multiple Label Learning (MIMLL), self-supervised learning, and unsupervised clustering.

8. The computer-implemented method of claim 6, wherein processing the plurality of training images to generate the machine learning prediction model further comprises:

receiving a plurality of predictions or a plurality of vectors of at least one feature from a weakly-supervised tile-level learning module for the plurality of training tiles;

training a machine learning model to take, as an input, the plurality of predictions or the plurality of vectors of the at least one feature from the weakly-supervised tile-level learning module for the plurality of training tiles; and

predicting a plurality of labels for a slide or a patient specimen, using the plurality of training tiles.

9. The computer-implemented method of claim 8, wherein at least one of the plurality of labels is binary, categorical, ordinal or real-valued.

10. The computer-implemented method of claim 8, wherein training the machine learning model to take, as the input, the plurality of predictions or the plurality of vectors of the at least one feature from the weakly-supervised tile-level learning module for the plurality of training tiles comprises a plurality of image features.

11. The computer-implemented method of claim 1, wherein the trained machine learning prediction model predicts at least one label using at least one unseen slide.

12. A system for processing electronic slide images corresponding to a tissue specimen, the system comprising:

at least one memory storing instructions; and

at least one processor configured to execute the instructions to perform operations comprising:

receiving one or more electronic slide images associated with a tissue specimen, the tissue specimen being associated with a patient and/or medical case;

partitioning a first slide image of the one or more electronic slide images into a plurality of tiles;

detecting a plurality of tissue regions of the first slide image and/or plurality of tiles to generate a tissue mask;

determining whether any of the plurality of tiles corresponds to non-tissue;

removing any of the plurality of tiles that are determined to be non-tissue;

determining a prediction, using a machine learning prediction model, for at least one label for the one or more electronic slide images, the machine learning prediction model having been generated by processing a plurality of training images; and

outputting the prediction of the trained machine learning prediction model.

13. The system of claim 12, wherein the plurality of tiles that are determined to be non-tissue are further determined to be a background of the tissue specimen.

14. The system of claim 12, wherein detecting the plurality of tissue regions comprises segmenting the tissue regions from a background of the one or more electronic slide images.

15. The system of claim 14, further comprising:
upon segmenting the tissue regions from the background, generating a tissue mask, the segmenting using thresholding based on color/intensity and/or texture features.

16. The system of claim 12, wherein the plurality of training images comprise a plurality of electronic slide images and a plurality of target labels.

17. The system of claim 12, wherein processing the plurality of training images comprises:

receiving a collection of digital images associated with at least one training tissue specimen, wherein the collection of digital images comprises a plurality of training electronic slide images;

receiving a plurality of synoptic annotations comprising one or more labels for each of the plurality of training electronic slide images;

partitioning one of the plurality of training electronic slide images into a plurality of training tiles for the plurality of training electronic slide images;

segmenting at least one tissue region from a background of the one or more electronic slide images to create a training tissue mask;

removing at least one of the plurality of tiles that are detected to be non-tissue; and

training the machine learning prediction model under weak supervision to infer at least one multi-label tile-level prediction using at least one label of the plurality of synoptic annotations.

18. The system of claim 17, wherein training the machine learning prediction model under weak supervision comprises using at least one of MIL, MIMLL, self-supervised learning, and unsupervised clustering.

19. The system of claim 17, wherein processing the plurality of training images to generate the machine learning prediction model further comprises:

receiving a plurality of predictions or a plurality of vectors of at least one feature from a weakly-supervised tile-level learning module for the plurality of training tiles;

training a machine learning model to take, as an input, the plurality of predictions or the plurality of vectors of the at least one feature from the weakly-supervised tile-level learning module for the plurality of training tiles; and

predicting a plurality of labels for a slide or a patient specimen, using the plurality of training tiles.

20. A non-transitory computer readable medium storing instructions that, when executed by a processor, cause the processor to perform a method for processing electronic slide images corresponding to a tissue specimen, the method comprising:

receiving one or more electronic slide images associated with a tissue specimen, the tissue specimen being associated with a patient and/or medical case;

partitioning a first slide image of the one or more electronic slide images into a plurality of tiles;

detecting a plurality of tissue regions of the first slide image and/or plurality of tiles to generate a tissue mask;

determining whether any of the plurality of tiles corresponds to non-tissue;

removing any of the plurality of tiles that are determined to be non-tissue;

determining a prediction, using a machine learning prediction model, for at least one label for the one or more electronic slide images, the machine learning prediction model having been generated by processing a plurality of training images;

and

outputting the prediction of the trained machine learning prediction model.