OBLIQUE DECISION TREES AS AN IMAGE MODEL FOR CUBIST IMAGE RESTYLING

Edric Chan
Great Oak High School

Magzhan Gabidolla   Miguel Á. Carreira-Perpiñán
Department of Computer Science and Engineering, University of California, Merced

IEEE International Conference on Image Processing (ICIP 2025)
14-17 September 2025, Anchorage, Alaska

ABSTRACT

We propose the use of oblique decision trees and forests, originally a supervised machine learning approach, as a model of a single image. This results in a hierarchical partition into constant-color polygons ("poxels") that best approximates the source image. The tree or forest can be trained using a recent algorithm that applies alternating optimization over the nodes of a tree and the trees of a forest. In this paper, we use this for aesthetic effects, such as restyling an image to look like a cubist painting or a stained-glass window. We illustrate this with multiple examples and compare it with the results achieved by Generative AI and Neural Style Transfer.

OBLIQUE TREES AND FORESTS AS AN IMAGE MODEL

- We propose to use oblique trees and forests as a model for an image.
- Decision trees have existed for a long time in ML, but traditional tree learning:
  - uses algorithms that are quite suboptimal (e.g. CART or C5.0)
  - is limited in the form of the trees (usually axis-aligned).
- We seek a tree-based model that can better approximate images, and an algorithm that can optimize it well.
- We use recent algorithms (TAO, FAO) that can learn a tree or forest for arbitrary choices of loss function and tree type (Carreira-Perpiñán et al. NeurIPS 2018, CVPR 2023).
- Our oblique tree/forest model has multiple uses (e.g. segmentation, compression, search), but here we focus on an aesthetic application: to restyle an image.
- This has applicability in software packages such as PhotoShop or for game and film production effects.

OBLIQUE TREES AND FORESTS AS AN IMAGE MODEL (cont.)

An oblique tree approximates a source image (considered as function from 2D to color space) by a hierarchical piecewise-constant function:

- each piece is a convex, colored polygon: poxel (for polygon pixel)
- its fidelity can be controlled by the model size (tree depth and number of trees).

The tree function defines the output image. It is a vector graphics object (not a bitmap).

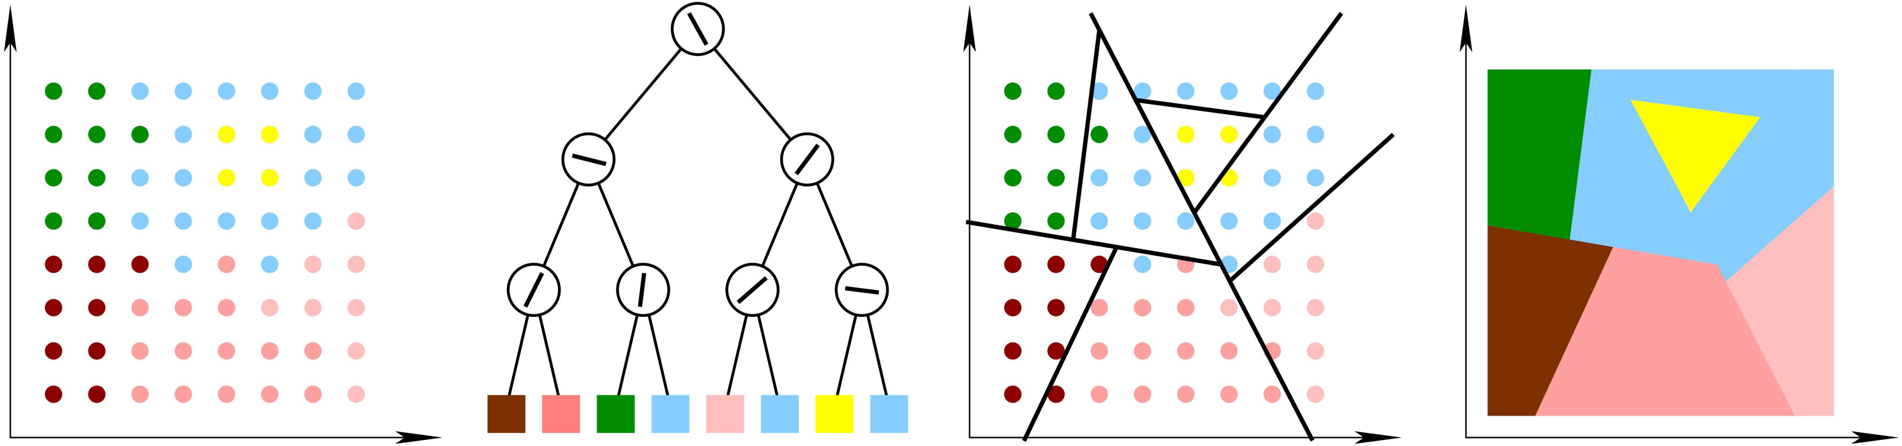

OBLIQUE TREES AND FORESTS AS AN IMAGE MODEL (cont.)

Some well-known paintings:

Mona Lisa (Leonardo da Vinci)

Starry Night (Vincent van Gogh)

Les Demoiselles d'Avignon (Pablo Picasso)

OBLIQUE TREES AND FORESTS AS AN IMAGE MODEL (cont.)

The output image from the oblique tree or forest is strikingly reminiscent of cubist paintings such as those of Picasso and other artists.

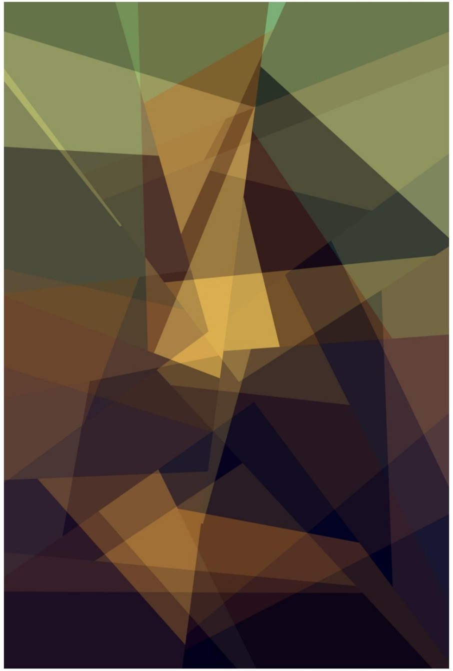

OBLIQUE TREES AND FORESTS AS AN IMAGE MODEL (cont.)

Other effects are possible, such as the style of stained-glass windows, as in Europe's gothic cathedrals. We obtain this by simply showing the poxel boundaries as thick lines.

Stained-glass window

Picture of Half Dome, Yosemite

Output image

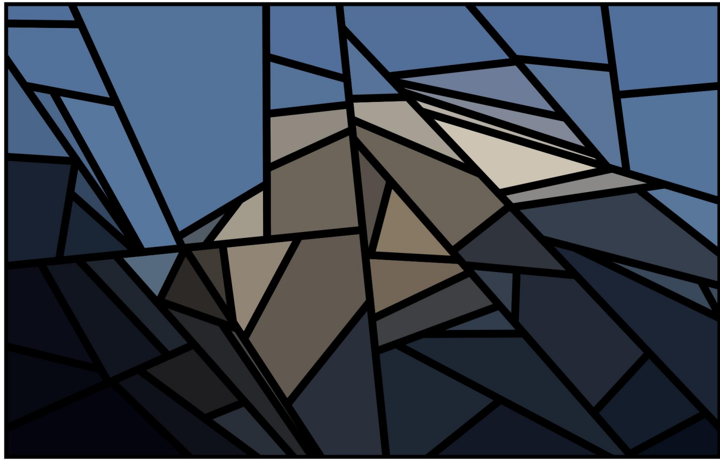

HOW TO LEARN THE TREE OR FOREST FROM AN IMAGE

We consider a regression problem where:

- the image is a labeled dataset of (location,color) for each pixel
- the regression function is a tree with parameters (split at decision nodes, label at leaves), or a forest (a sum of trees)
- we minimize the squared error of the image wrt the function.

To solve this we use these algorithms:

- Tree Alternating Optimization (TAO)
- Forest Alternating Optimization (FAO)

Both of them take an initial tree or forest and iteratively update its parameters so the objective function decreases until convergence.

TAO AND FAO

In Tree Alternating Optimization (TAO) we update each node in sequence:

- Separability condition: nodes at the same depth can be updated in parallel.
- Reduced problem over a decision node: updating the split is equivalent to a weighted 0/1 loss binary classification problem, where each pixel in the node tries to choose the child that best predicts it.
- Reduced problem over a leaf: the optimal leaf label is the average of the pixels it contains.

In Forest Alternating Optimization (FAO) we:

- update each tree in sequence using TAO
- update the labels of all the leaves in all the trees jointly

For details, see the references in the paper.

USER PARAMETERS AND CONTROL: EFFECTS

The following parameters can be tuned to achieve a variety of aesthetic effects:

- Tree depth Δ and number of trees T. They control the number of poxels:
  - Few poxels: cartoon-like.
  - The more poxels, the more faithful the approximation to the source image.
  - Many poxels: photorealistic.

  Using multiple trees gives an appearance of "broken glass".
- Seed for the initial tree (random median partitions). Useful to add some randomness.
- Line width of poxel boundaries. Useful to produce a "stained-glass windows" effect.
- Axis-aligned tree: a special case of an oblique tree where the split is either horizontal or vertical. It generates rectangular poxels (blocky Mondrian-like effect).
- Use a mask so that the tree or forest is learned on only the pixels within the mask.

TRAINING COMPLEXITY AND RUNTIME

TAO/FAO's complexity is linear on the number of pixels. Processing an image takes from less than a second to a few seconds in a laptop (depending on the parameters).

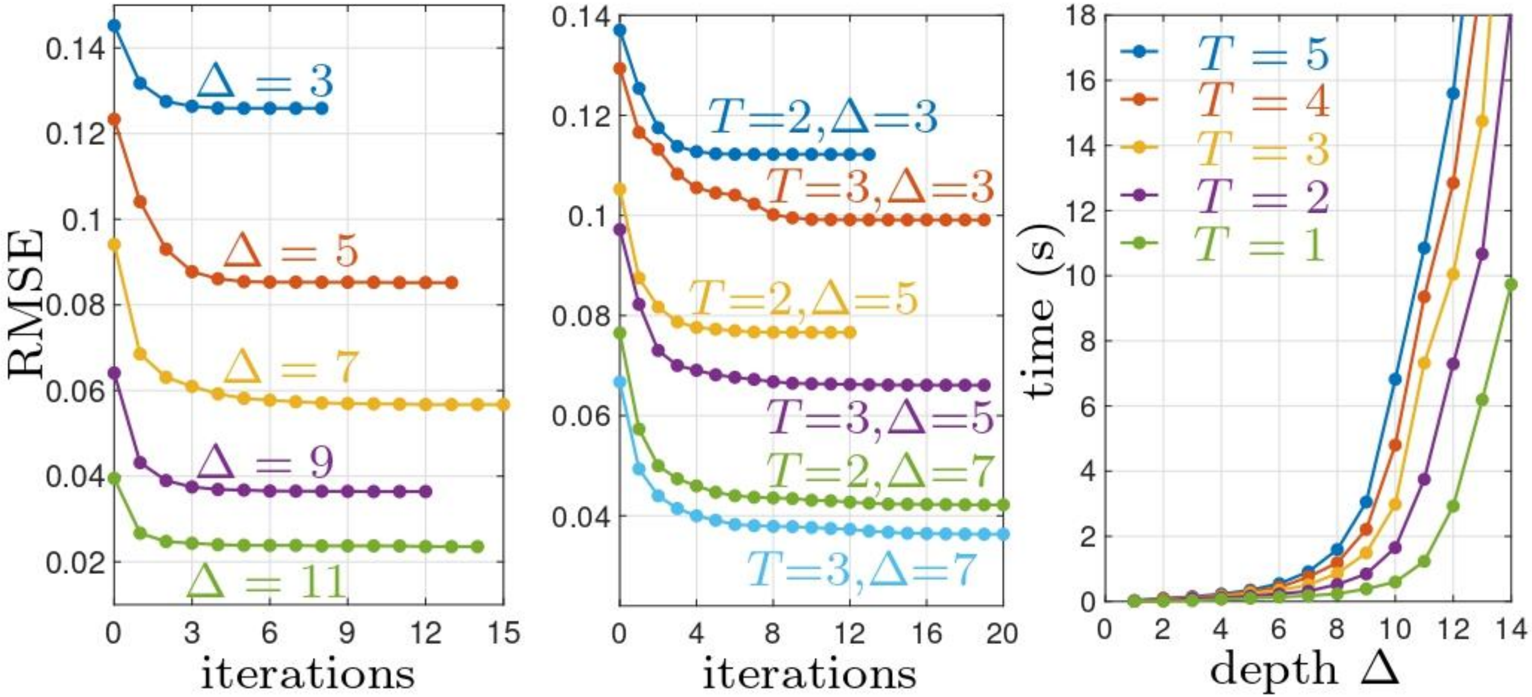

INITIALIZATION AND FINAL RESULT

Initial (random median trees) and final result (after TAO/FAO converge), as poxels and boundaries, for two models (with different parameters). Source image: Leonardo's *Mona Lisa*.

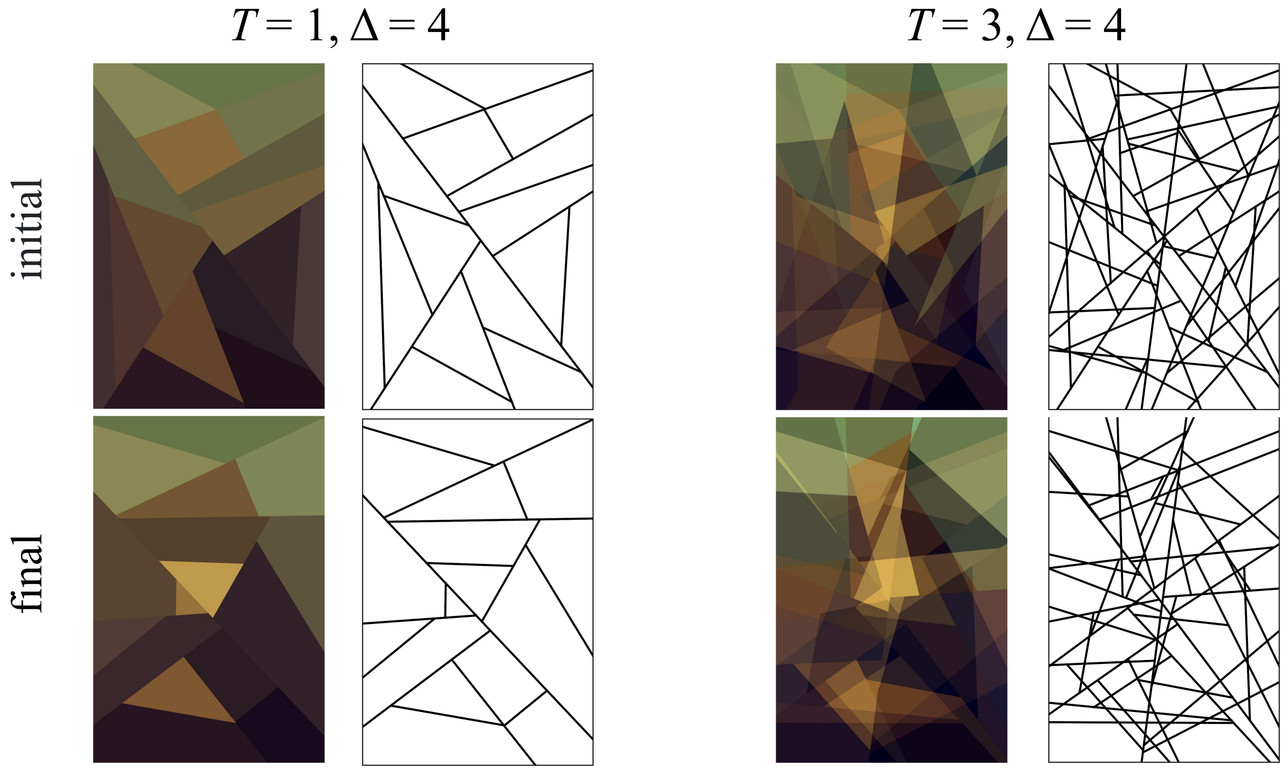

DIFFERENT DEPTH AND NUMBER OF TREES

source image

$T = 1, \Delta = 1$

$T = 1, \Delta = 5$

$T = 1, \Delta = 12$

$T = 2, \Delta = 7$

$T = 5, \Delta = 6$

DIFFERENT RANDOM SEEDS

Combining cubist tree images using different seeds into a video produces a jittery effect reminiscent of rotoscopic animation.

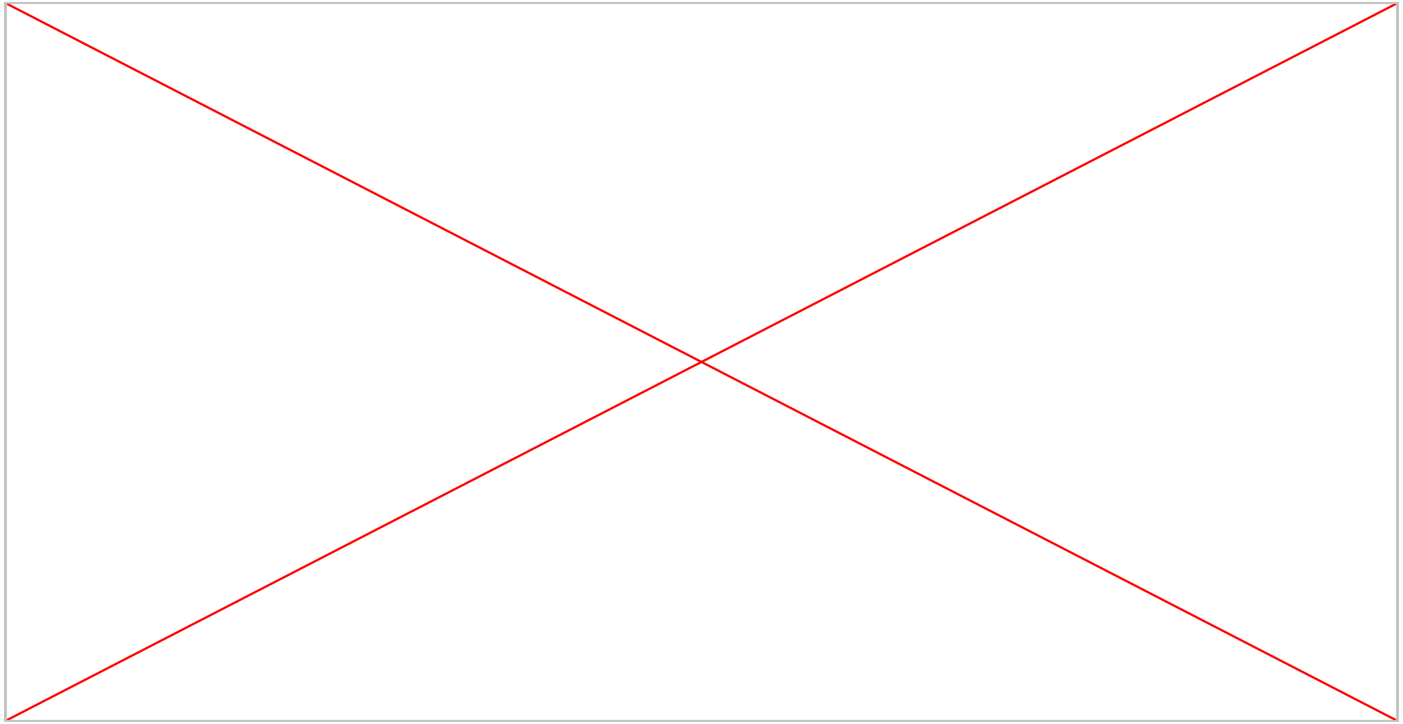

INDIVIDUAL TREES OF A FOREST

$T = 3$, $\Delta = 4$. Source image: Van Gogh's *Starry Night*.

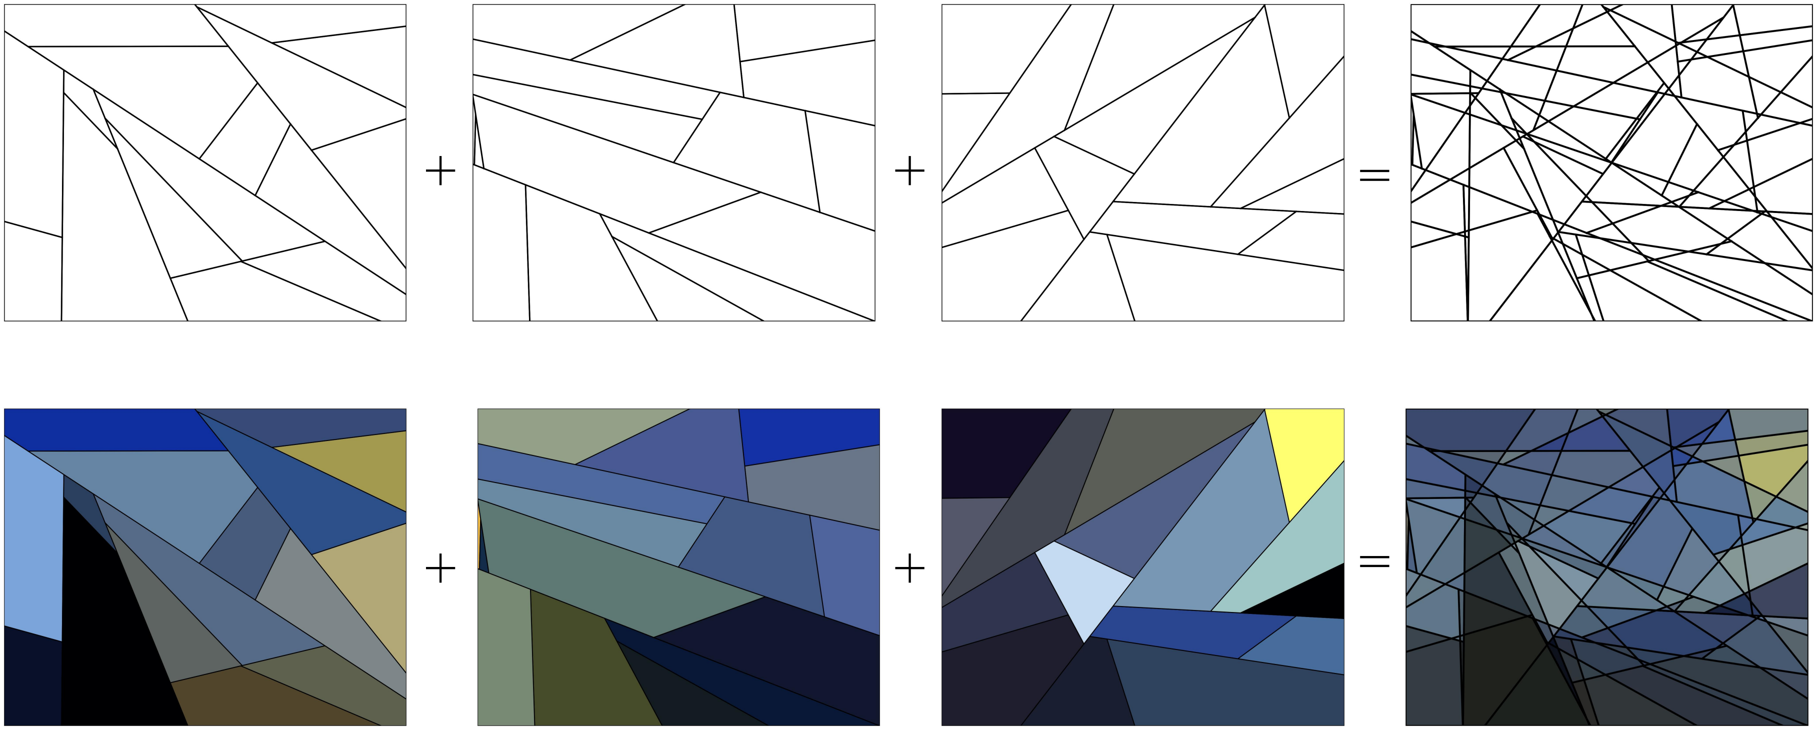

POSTPROCESSING A CUBIST IMAGE WITH OTHER SOFTWARE

The poxels are vector graphics objects and can be used by traditional image manipulation software. Source image: *Marilyn Monroe* by Andy Warhol.

Source image

Cubist tree output

Cubist tree output, edited

ART GALLERY

In the remaining slides, we show different cubist tree images and the tree/forest parameters. Can you guess the original paintings, drawings, cartoons or photographs?