A picture is worth 1000.. features!

Using computer vision alongside machine learning in computer systems.

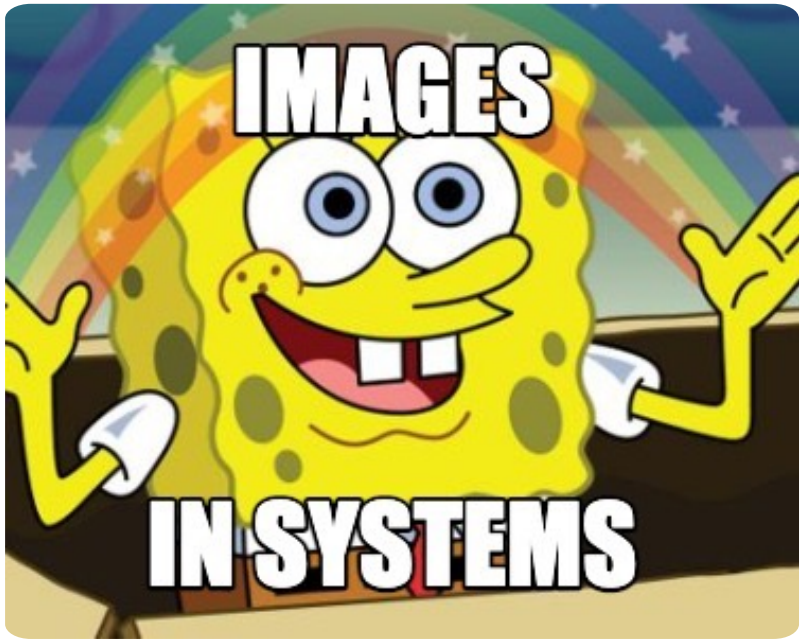

Thaleia Dimitra Doudali

Assistant Professor

IMDEA Software Institute

ASPLOS 2022 WACI, Lausanne, Switzerland

Relationship of Computer Systems Researchers with Visualization

My paper "**Kleio**" at HPDC 2019.

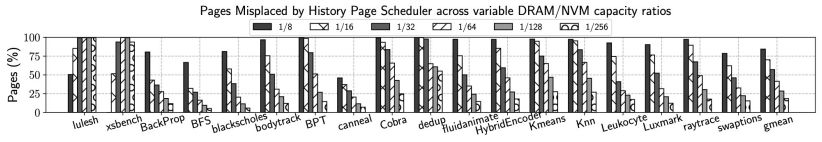

Figure 6: Percentage of pages misplaced at least one time across the scheduling epochs by the History page scheduler. This is the set of pages that need machine intelligence based management. This observation is crucial since it highlights that the problem space of per page RNN training can be significantly reduced as the size of available DRAM does.

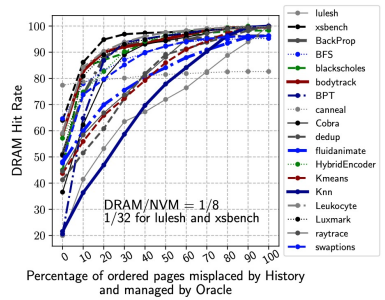

Figure 7: DRAM hit rate when an Oracle Page scheduler manages the misplaced-by-History pages and the History page scheduler manages the rest. Pages are ordered in descending performance benefit. Clever management of even a small percentage of these pages, can give most of the performance benefits we would have by managing cleverly all pages.

5.1 Page Selector

We first describe the page selector component in Kleio. Its design is driven by the following observations regarding the importance of correct page placement to application performance:

- There is only a certain subset of pages that needs more clever data management, than what the existing history-based solutions can provide. That subset is significantly small for limited DRAM capacity.
- Pages that need machine intelligence based management, can be ordered with respect to the performance impact of their placement into the appropriate memory component. We define a benefit metric that enables the page ordering, prioritizing pages with high access counts and number of misplacements by the History page scheduler.
- Intelligent management of the pages following the aforementioned ordering does not correspond to linear performance improvement. In contrast, intelligent placement for only (a small) part of them can bring most of the performance benefits we would get by applying intelligent placement across all application pages.

We define a 'misplacement' of a page by the History scheduler, when at the start of a scheduling epoch, a page was supposed to be allocated in DRAM, but it was not, because of wrong hotness prediction. Figure 6 depicts the percentage of application pages, which

hackercombat.com

We **make** graphs for system and paper evaluation.

We **look** at graphs to understand system behaviors.

My Relationship with Visualization

PhD

Topic: Machine Learning for Memory Management Systems*

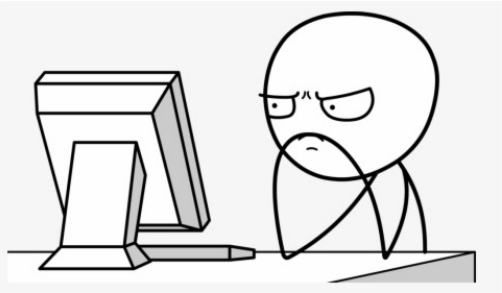

Spent years staring at these images.

I **visualized** "memory access patterns" to explain system behaviors across application domains.

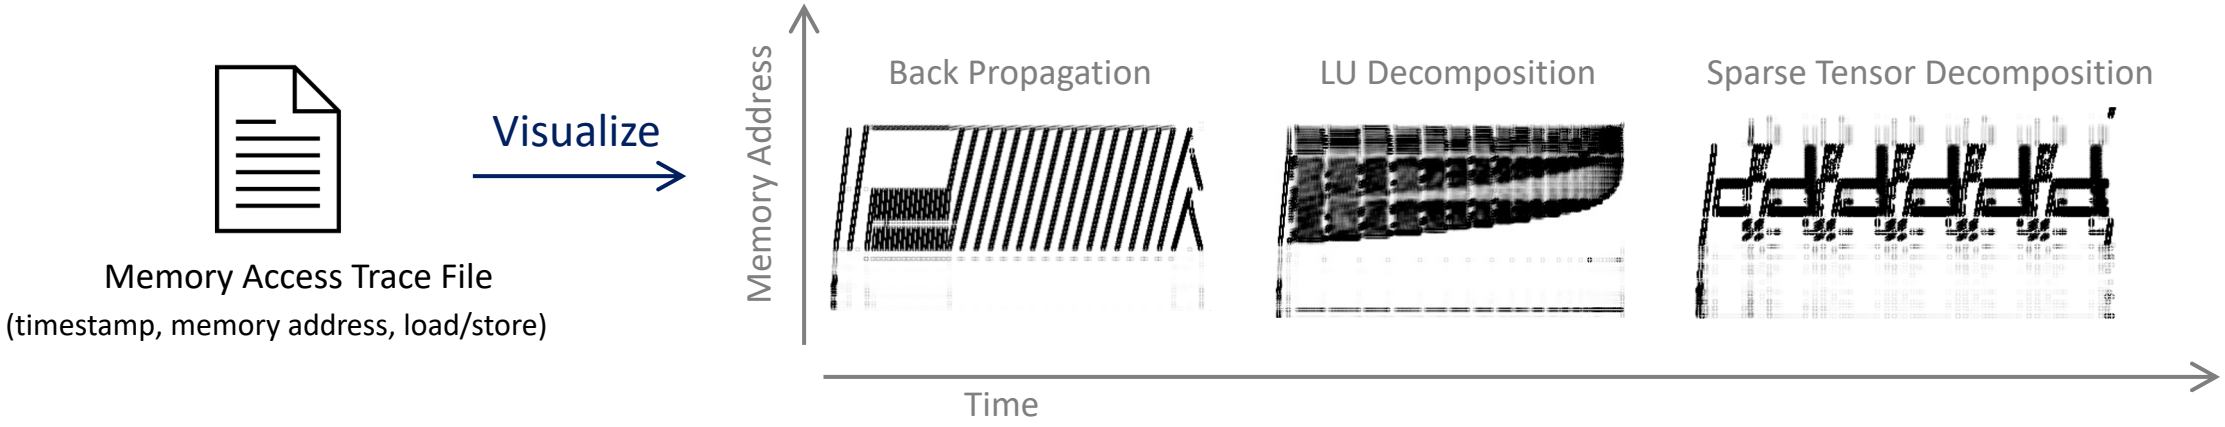

The key insights of my systems designs came from *visual* observations!

My papers: "**Cori**" at IPDPS 2022 and "**Coeus**" at CCGrid 2022.

*Adding Machine Intelligence to Hybrid Memory Management. Thaleia Dimitra Doudali. PhD Dissertation, Georgia Tech, 2021.

Visualizing Data Access Patterns

Let's create images for larger workloads.

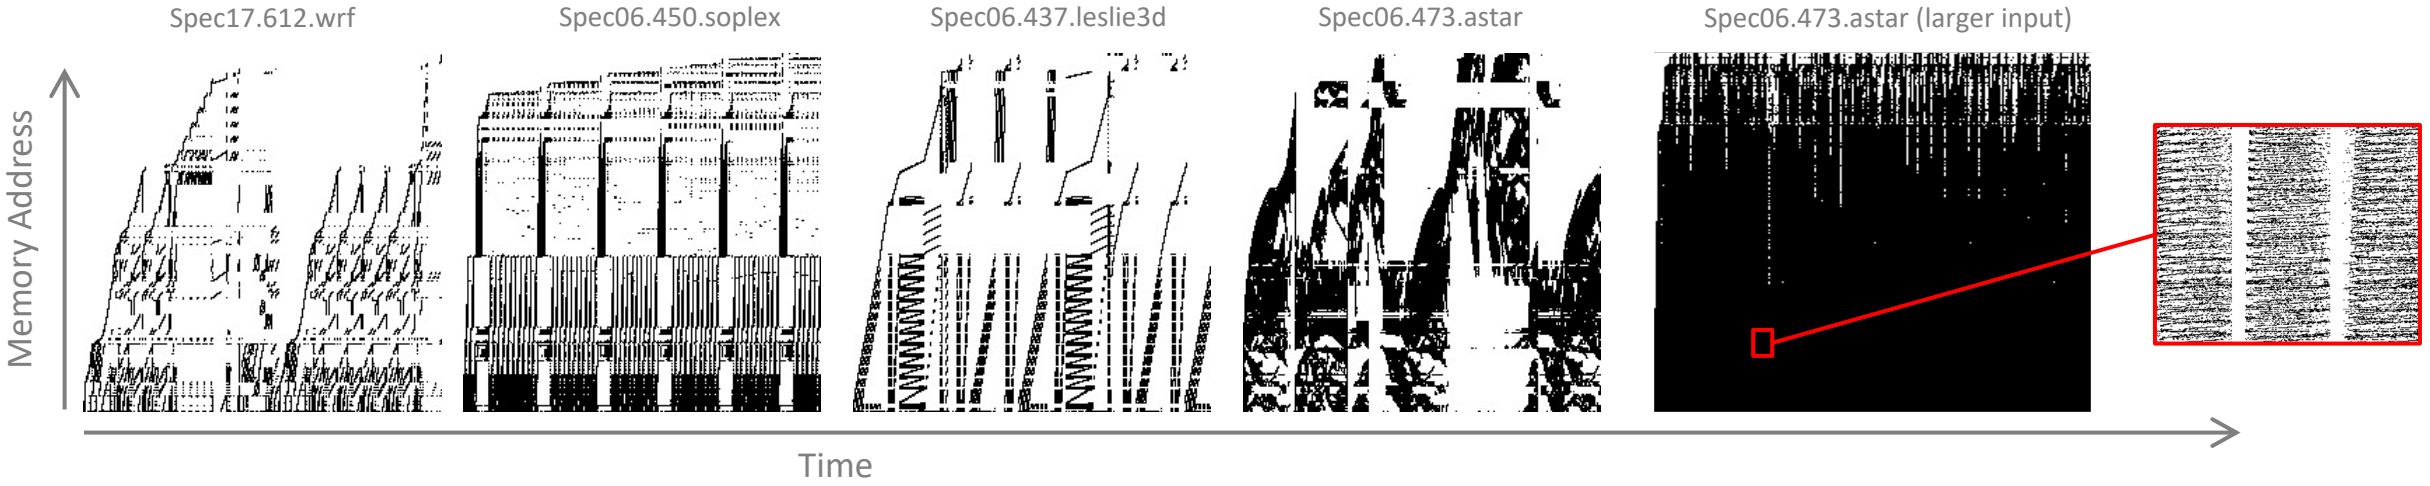

Challenge: limited 2D space to depict millions of data points.

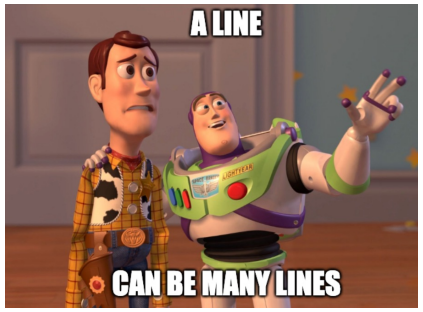

Let's zoom-in until we see clear lines.

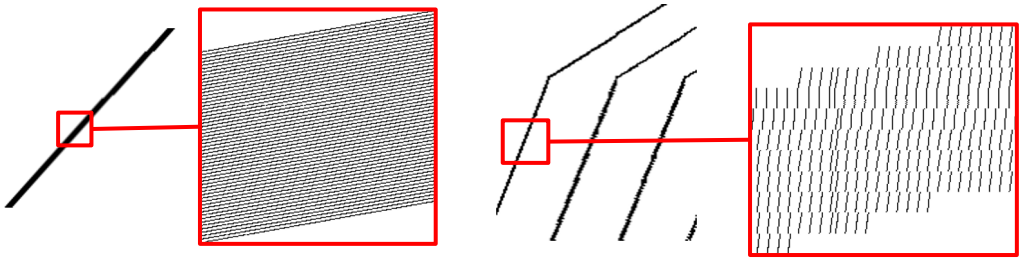

Open Problems:

- Make 1 image and then zoom in?
- Make many images from start?
- Time window per image?
- Image resolution / size / color?
- Metadata?
  Benchmark, level of mem/cache, etc..

Computer Vision + Machine Learning for Systems

Learning data access patterns.

What can an image-based system pipeline look like?

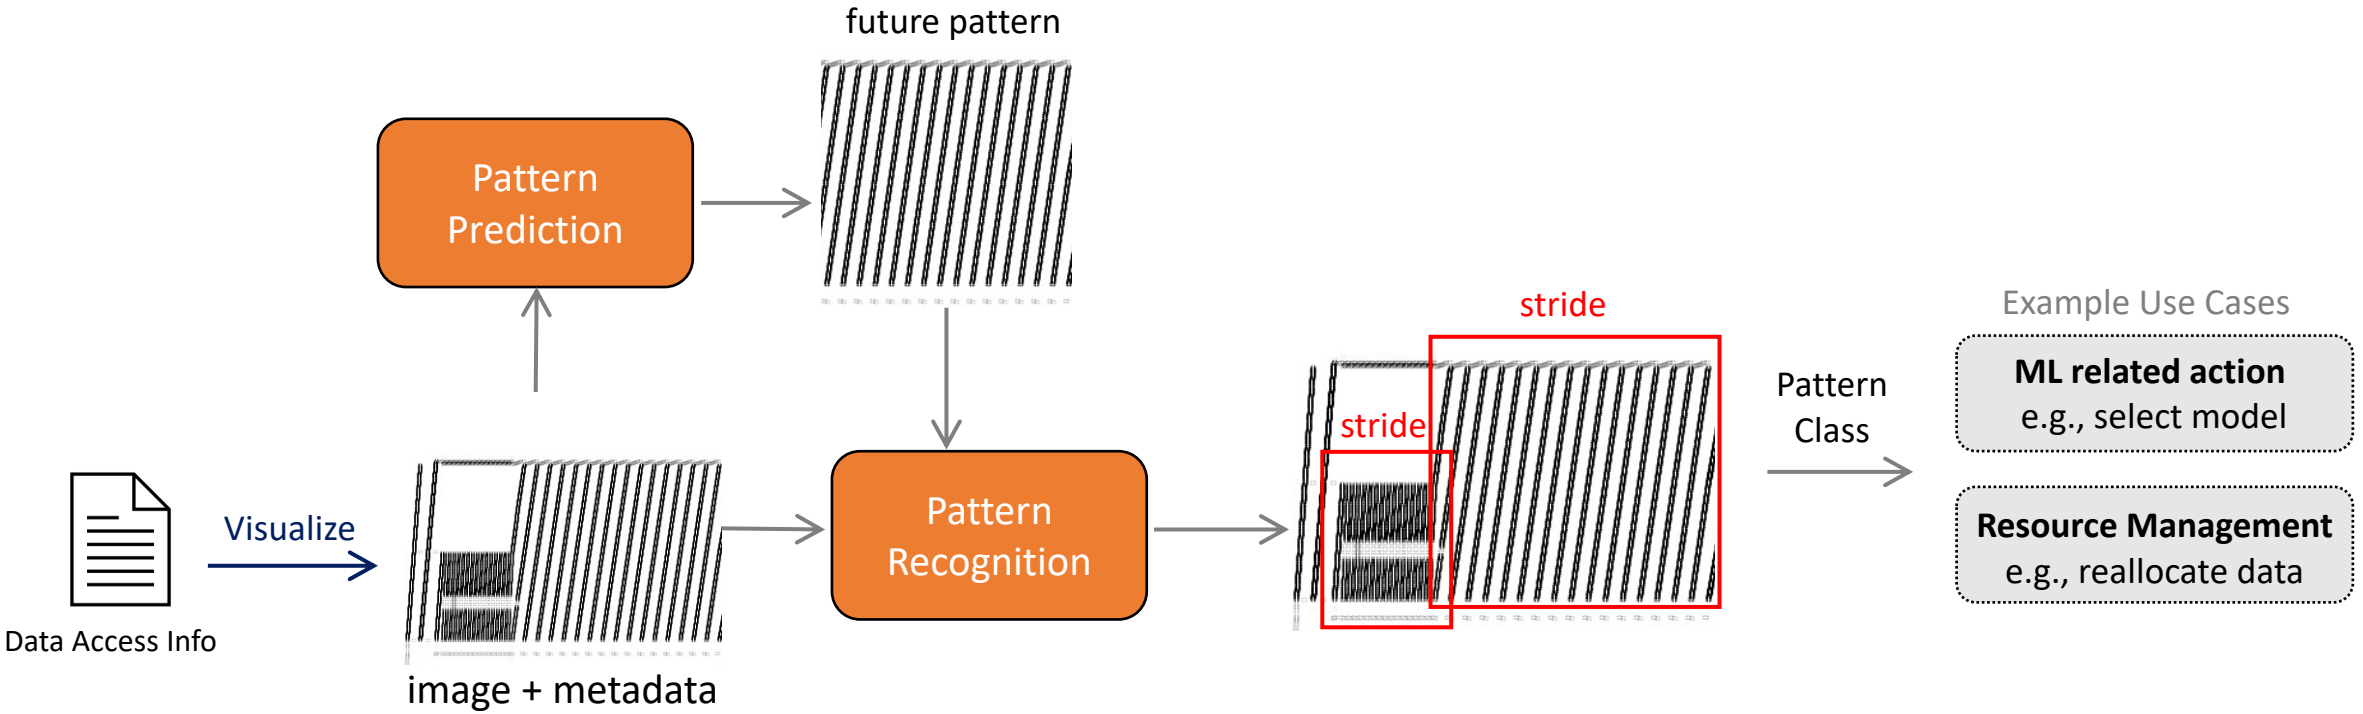

Future work building on: Adding Machine Intelligence to Hybrid Memory Management. Thaleia Dimitra Doudali. PhD Dissertation, Georgia Tech, 2021.

Pattern Recognition

After solving the visualization challenges described above..

..Can we build an “ImageNet”, a public image dataset of data access patterns?

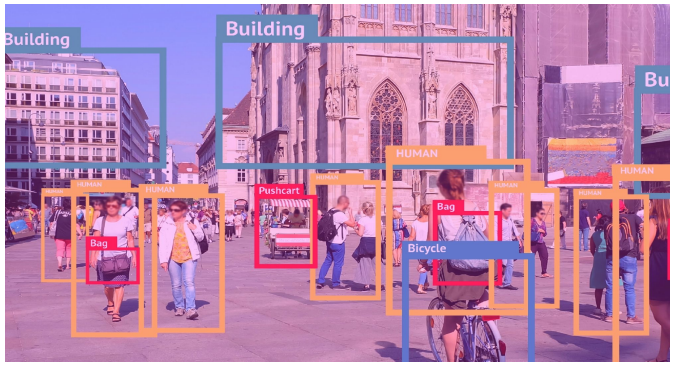

Can we create something similar?

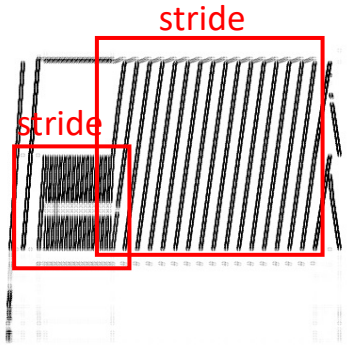

Challenge: how to properly label data access patterns?

Open Problems:

- What classes to define?
- Labeling guidelines?
- Community Contributions?
- Train classifiers for pattern recognition?
- Impact of misclassification?
- OS/Library/Compiler/Runtime support for pattern detection?

Pattern Prediction

As the workload is running, the Operating System (OS) is “watching” a **video** of how the application accesses data.

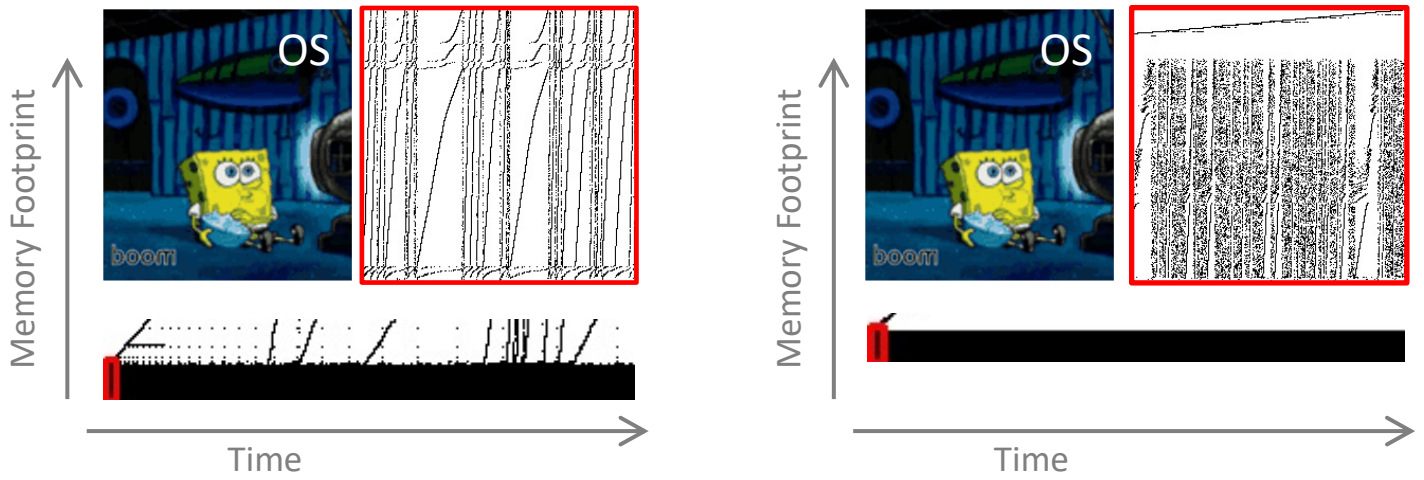

Rolling window across time for part of the memory footprint.

Open Problems:

- Which one is “best”:
  - CV + ML vs. ML vs. non ML.
  - Accuracy, training times, misprediction impact.
- 1 model for all, per app, per pattern?
- Training intervals vs. OS operation.

We can leverage machine learning methods for predicting the next frame of a video.

Other Use Cases of Computer Vision in Operating Systems

Forecasting any type of time series data: e.g., server / cloud / application resource usage over time.

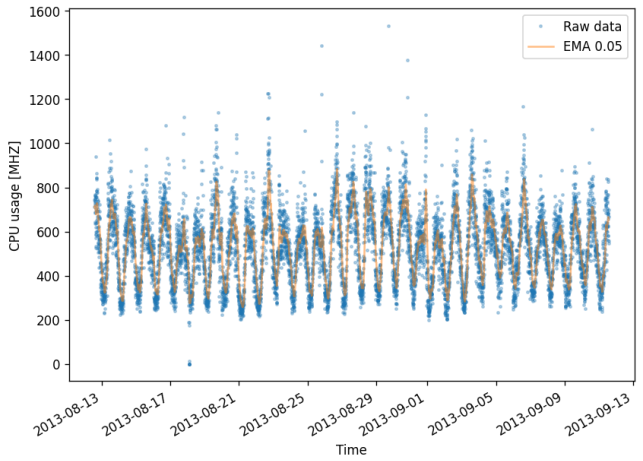

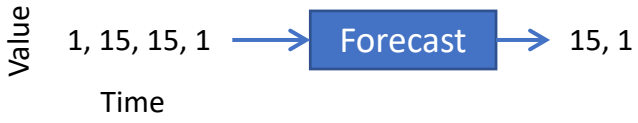

Learning raw values

VS.

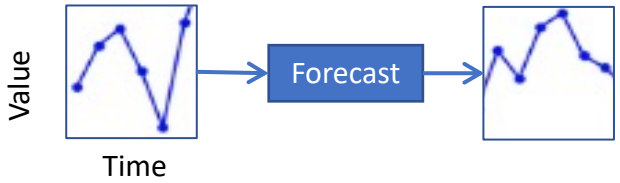

Learning images

Which one will be "better"?

In the financial domain learning images lead to higher accuracy.

From OSDI '21 Keynote from J.P. Morgan AI labs.

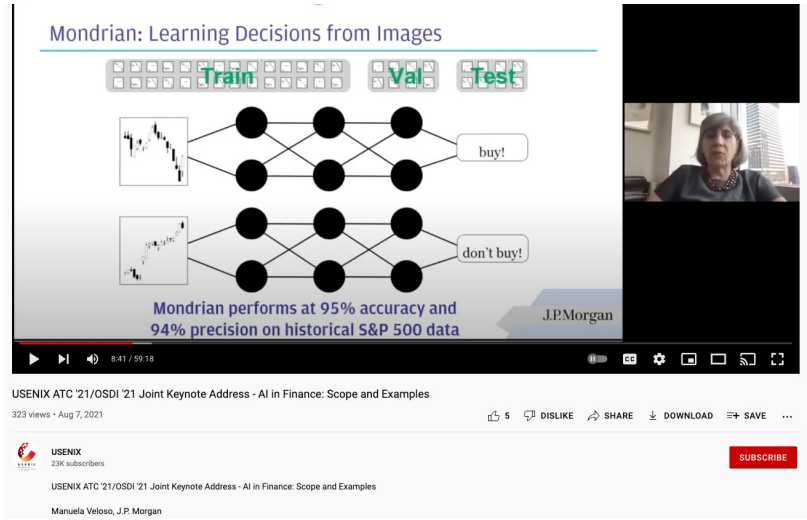

Open Problems:

- How many timesteps per image?
- Similar visualization and labeling challenges as above.

Why Images?

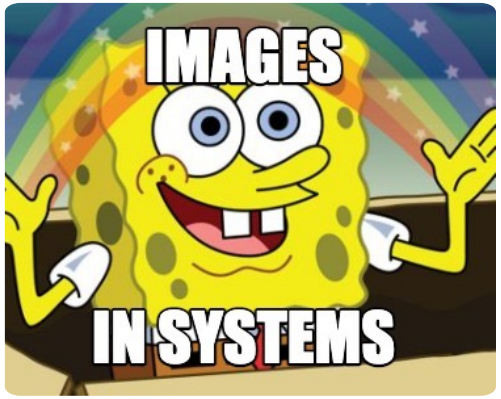

Let’s rethink how we represent data for machine learning.

Creating images helps:

- Reduces dimensionality to a 2D space (3D if color). **A picture is worth 1000.. Features!**
- Captures spatial and temporal correlations.
- Reduces input space and training times e.g., 10000 raw values vs. 10x10 image.
- Leverage computer vision algorithms.

Will it be more effective than ML or non ML solutions? Let’s see!

Just the insight of observing those images is beneficial for how we design systems.

SysMLCV

Systems

+

Machine Learning

+

Computer Vision

I am proposing a new intersection of research areas, the **SysMLCV**.

Let’s work together!

Scan my website