Formalizing Interruptible Algorithms for Human *over*-the-loop Analytics

Austin P. Graham
University of Oklahoma
austin.p.graham-1@ou.edu

The University of Oklahoma

Key Problems

Learning algorithms are completely autonomous.

Predictive models are hard to digest.

Humans are slow thinkers.

 (Graham et al., 2017)

$$
\begin{aligned}
&(x-2)^2(y-2x+2)^2(y+2x-10)^2(x-4)^2(y-2x+8)^2(y+2x-16)^2\left(y-3-3\left\lfloor x-\tfrac{11}{2}\right\rfloor^2\right)^2(x-8)^2\\
&\cdot\left(y-2-3\left\lfloor \frac{x-8}{2}\right\rfloor^2\right)^2(x-11)^2\left(y-\tfrac{1}{2}x+\tfrac{5}{2}-3\left\lfloor \frac{x-11}{2}\right\rfloor^2\right)^2\left(y+\tfrac{1}{2}x-\tfrac{17}{2}-3\left\lfloor \frac{x-11}{2}\right\rfloor^2\right)^2(x-15)^2\\
&\cdot\left(y-4-3\left\lfloor \frac{x-14}{2}\right\rfloor^2\right)^2(y-2x+52)^2(x-17)^2(y+x-21)^2(x-19)^2\left(y-x+17-3\lfloor x-20\rfloor^2\right)^2\\
&\cdot\left(y+x-23-3\lfloor x-20\rfloor^2\right)^2\left(y-x+19-3\lfloor x-21\rfloor^2\right)^2\left(y-3-3\lfloor x-21\rfloor^2\right)^2(x-25)^2\left(y+\tfrac{1}{4}x-\tfrac{41}{4}-3\left\lfloor \frac{x-25}{2}\right\rfloor^2\right)^2\\
&\cdot\left(y-\tfrac{1}{8}x-\tfrac{1}{8}-3\left\lfloor \frac{x-25}{2}\right\rfloor^2\right)^2\left(y+\tfrac{5}{8}x-\tfrac{151}{8}-3\left\lfloor \frac{x-25}{2}\right\rfloor^2\right)^2(y-2x+54)^2(y+2x-62)^2\left(y-3-3\left\lfloor x-\tfrac{57}{2}\right\rfloor^2\right)^2\\
&\cdot(x-31)^2(y+x-35)^2(x-33)^2(x-34)^2\left(y+\tfrac{1}{2}x-21-3\left\lfloor \frac{x-34}{2}\right\rfloor^2\right)^2\left(y-\tfrac{1}{2}x+15-3\left\lfloor \frac{x-34}{2}\right\rfloor^2\right)^2\\
&\cdot((x-38)^2+(y-3)^2-1)^2(x-40)^2(y+2x-84)^2(y-2x+80)^2(x-42)^2(x-43)^2\left(y-2-3\left\lfloor \frac{x-43}{2}\right\rfloor^2\right)^2\\
&\cdot(y-3-|x-47|)^2((x-47)^2+(y-3+\sqrt{y^2-6y+9})^2)^2+\left(y^2-6y+8+\sqrt{y^4-12y^3+52y^2-96y+64}\right)^2=0
\end{aligned}
$$

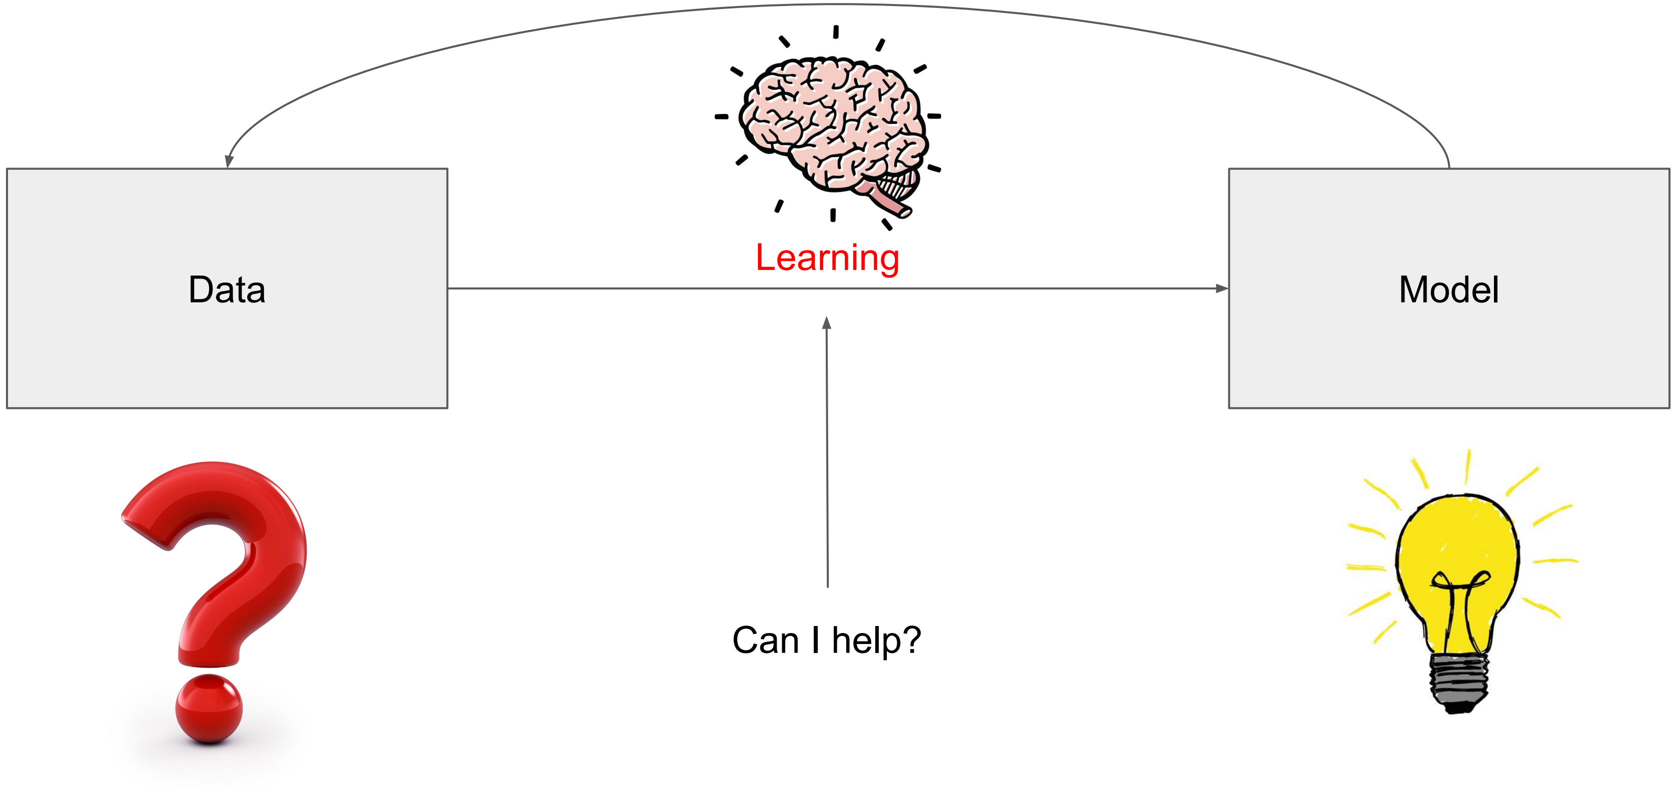

 (Graham et al., 2017)

Anytime Algorithms

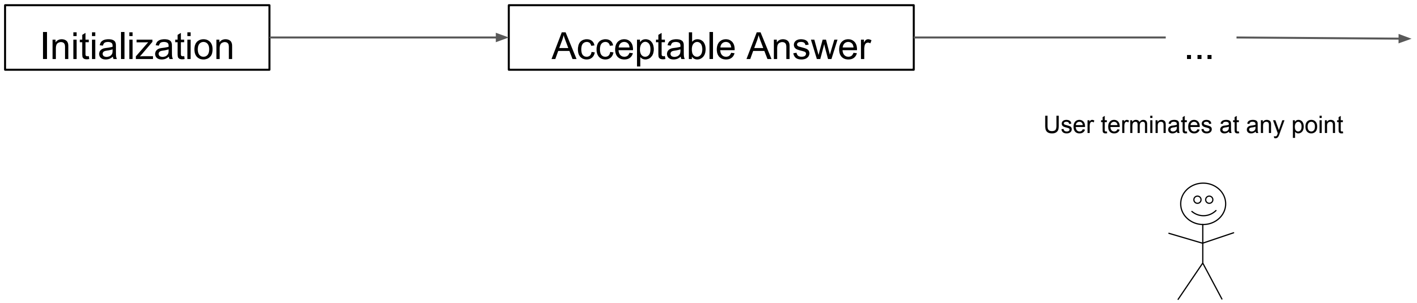

Interactive Algorithms

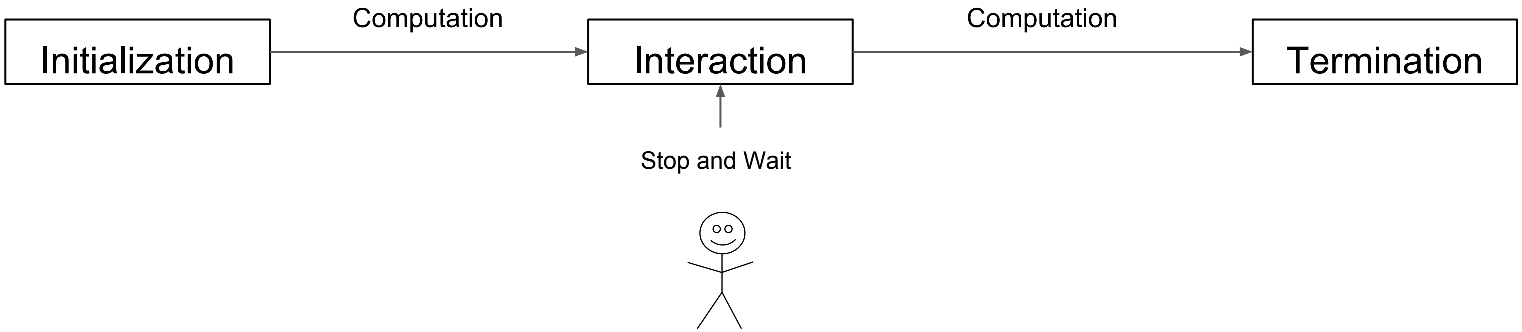

Interruptible Algorithms

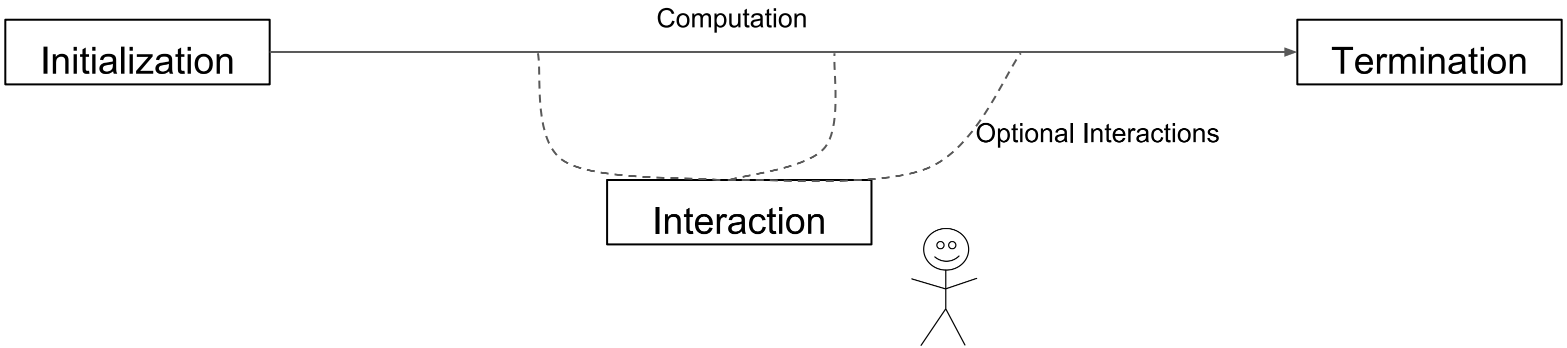

How can I help?

Repose the question → Change model attributes

Give more information → Add/Remove data

 (Graham et al., 2017)

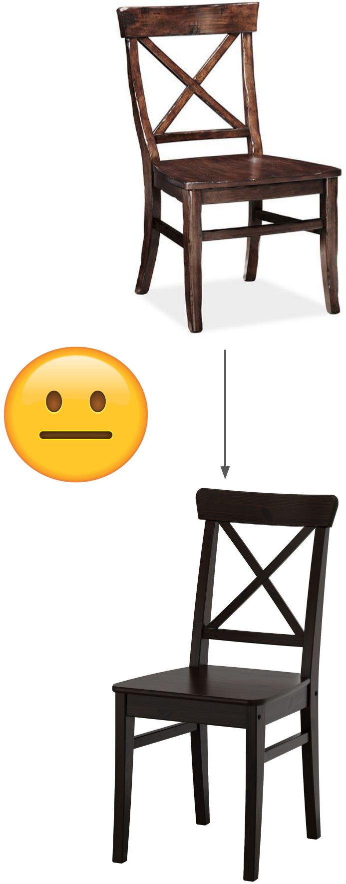

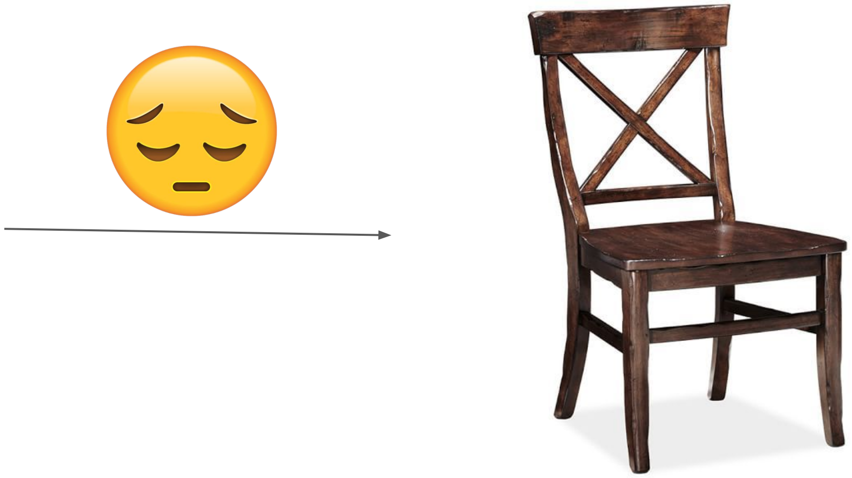

Key phrase: might as well start over

Minimally Increase Execution Time for Better Results

$$I_{A,U} > O(A)/O(C)$$

How hard the algorithm is

How hard the change is

Greater than 1:
Okay to interrupt

Less than 1:
Just start over

 (Graham et al., 2017)

Key Questions

What kind of changes can be made to positively affect outcomes?

How do I know I’ve made things worse?

 (Graham et al., 2017)

Visualization

Is there a model agnostic way to visualize learning?

Will the visualizations be readable to non-experts?

In what contexts will humans have enough time to react?

 (Graham et al., 2017)

Questions?

Motivations - Data Mining and Machine Learning

- Long running
  - Days, Weeks
  - Changes in parameters must be made *a posteriori*
- Static
  - Changing data means re-training model

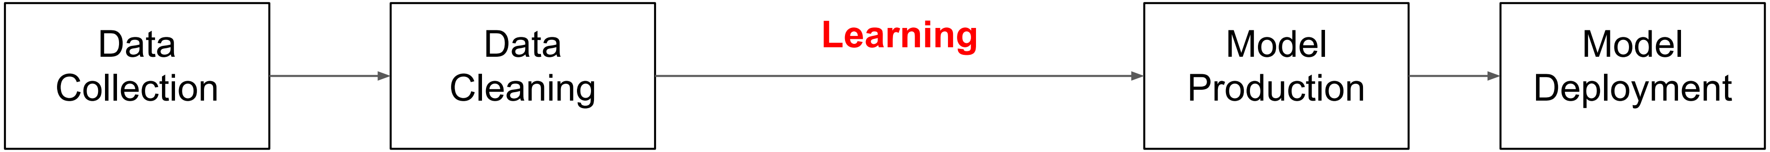

 (Graham et al., 2017)

Humans Getting Involved

- Interactive
  - Humans *IN*-the-loop
  - Focus on model accuracy
- Anytime
  - Humans *ENDING*-the-loop
  - Focus on model availability
- Interruptible
  - Humans *OVER*-the-loop
  - Manage the accuracy/availability tradeoff

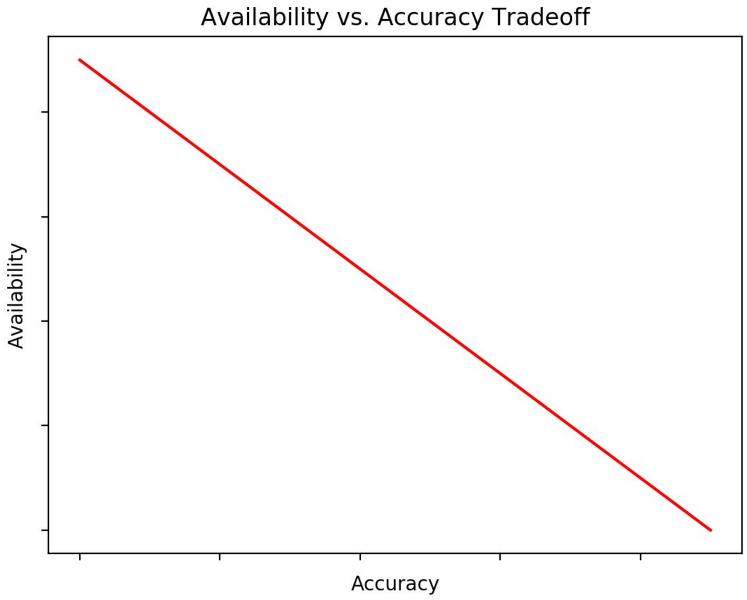

 (Graham et al., 2017)

Interactive Algorithms

- Learning models asking humans questions
  - Stop-and-wait conditions
- Improves accuracy with detrimental increases to runtime

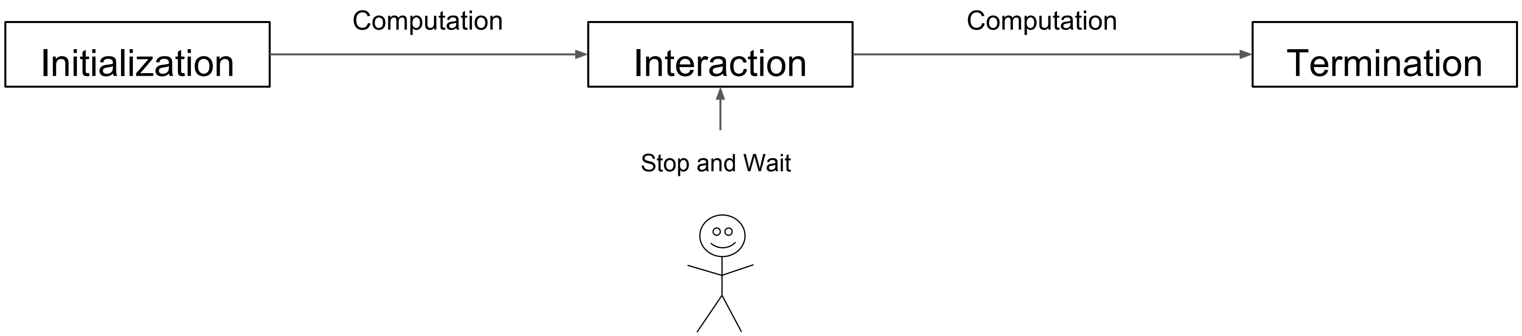

Example Interactive Approaches

- Awasthi et al.
  - Hierarchical clustering
  - Analyst can split/merge clusters at every level
- Lad and Parikh
  - Image clustering
  - Algorithm asks the analyst for the answer
- Amershi et al.
  - Clustering in social networks
  - Improvement happens on user choice

Anytime Algorithms

- Learning models build to an acceptable point, then improve until user says to stop or convergence
  - Assumes models improve with longer runtimes
- Allows analysts to train their comfort level

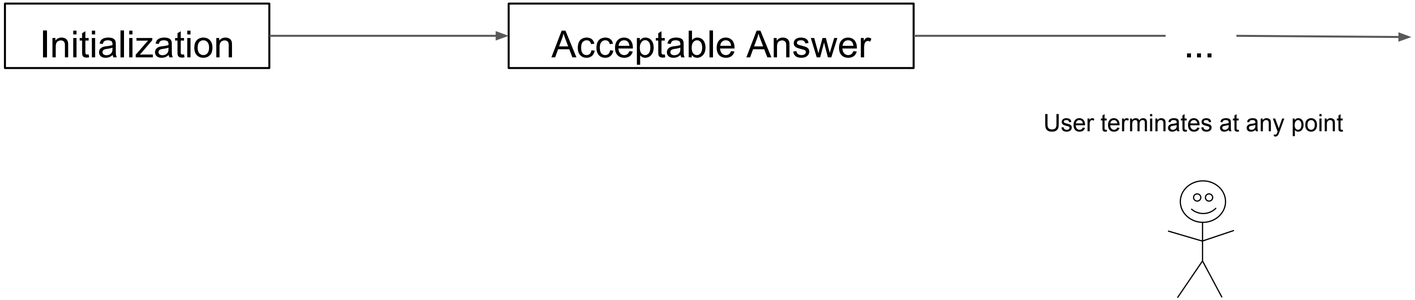

 (Graham et al., 2017)

Anytime Approaches

- Ueno et al.
  - Stream mining
  - Anytime structure accounts for fluctuations in data streams
- Vlachos et al.
  - Time Series clustering
  - Coarse clustering for a rough estimate, then iteratively improve.

 (Graham et al., 2017)

Meet in the Middle - Interruptible Algorithms

- Learning algorithms follow predefined behavior allowing analyst to make changes at available times
  - Analyst can intervene only if they find it necessary
- Increases in execution time depend on user involvement

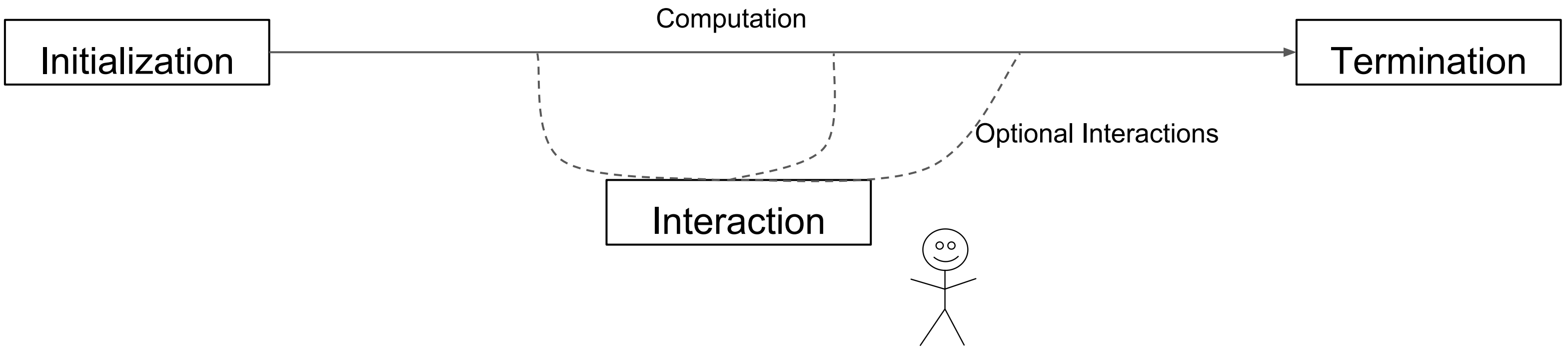

The Interruptibility Index

- Measures change affects relative to the algorithm

$$I_{A,U} > O(A)/O(C)$$

- “Algorithm *A* is interruptible in attribute *U* if the ratio of the complexity of *A* to the complexity of a change *C* in *U* is > 1”

$$I_{A,U} > I_{A,S}$$

- “If the interruptible index in attribute *U* for algorithm *A* is greater than the interruptibility index in attribute *S*, then the change in *U* is less interruptible than the change in *S*”

 (Graham et al., 2017)

HOLA - An Interruptible Algorithm Provider

- (H)uman (O)ver the (L)oop (A)nalytics
- Architecture
  - User Space
    - Interact with an analyst
    - Visualization system conveying model state
    - Convert user gestures and actions to machine readable 'system calls'
  - Kernel Space
    - Manage concurrent changes in data and model information
    - Store and manage raw data set
    - Return renderable model state to User Space

 (Graham et al., 2017)

System Calls

- Atomic operations that make up interrupts
- Data Changes
  - Update a record
  - Remove a record
  - Add a record
- Data Model Changes
  - Update hyper parameter
  - Inspect model state
- Control Changes
  - Open connection
  - Close connection

 (Graham et al., 2017)

Architecture

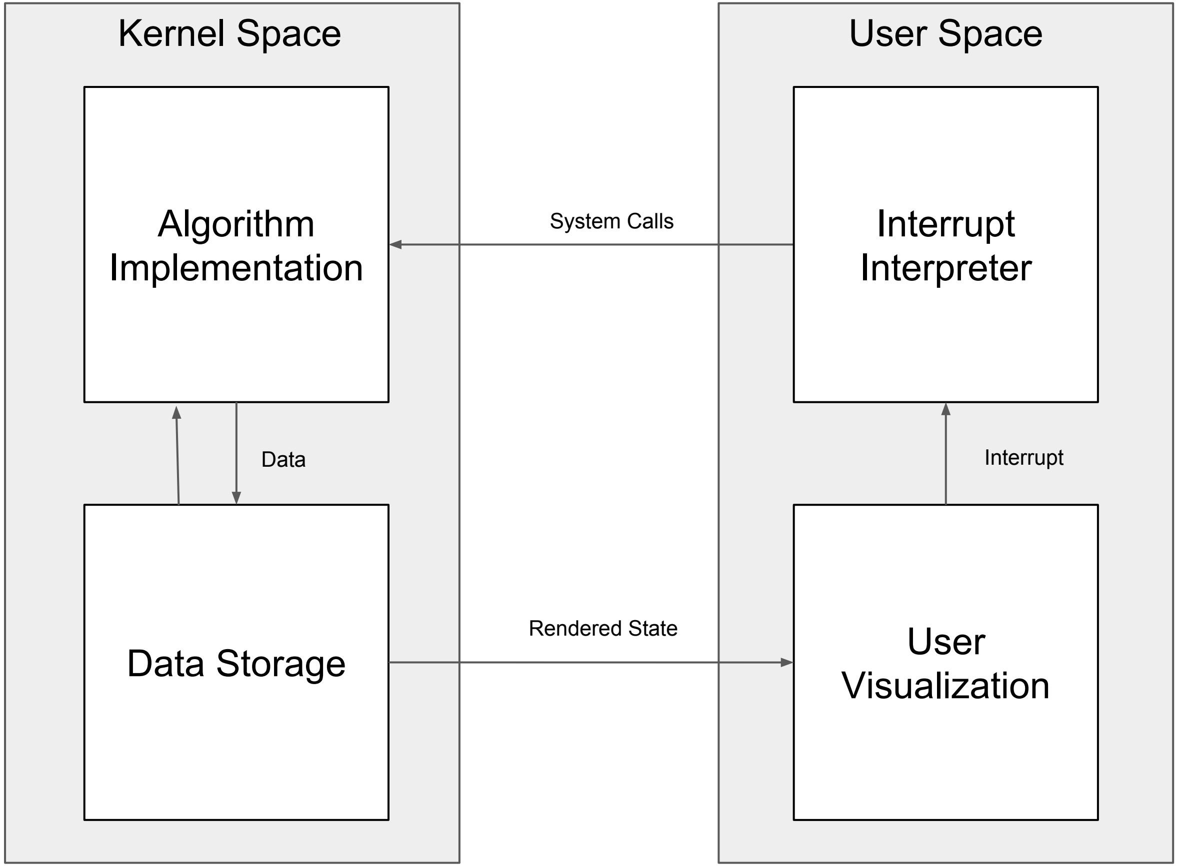

 (Graham et al., 2017)

Interruptibility Example: KMeans Clustering

- KMeans Algorithm:
  - Initialize k clusters
  - Randomly assign points to each cluster
  - Calculate initial cluster centers
  - While a change has been made
    - Calculate closes cluster center to each point
    - Reassign the point to closest cluster
- Decisions to be made
  - What changes are interruptible?
  - Where should the interrupts happen?

 (Graham et al., 2017)

KMeans Interruptibility Indices

- Complexity of KMeans = $O(nkI\theta)$
- Complexity of changing a single data point *q* used by KMeans = $O(\theta)$
- $I_{KMeans,q} = O(nkI\theta)/O(\theta) = nkI$
  - *n*, *k*, *I* are constants with respect to a data point
  - Interrupt is constant time
- Complexity of changing k hyper parameter in KMeans = $O(\theta)$
- $I_{KMeans,q} = O(nkI\theta)/O(\theta) = nkI$
  - *n*, *I* are constants with respect to k hyper parameter
  - Interrupt is linear time

Where to Interrupt

Listing 1 Sample KMeans Interruptible Implementation

```
@hola(data_points, K)
def find_centers(data_points, K):
    oldmu = random.sample(X, K)
    mu = random.sample(X, K)
    while not has_converged(mu, oldmu):
        oldmu = mu
        hola.interrupt()
        clusters = cluster_points(X, mu)
        mu = reeval_centers(oldmu, clusters)
    return(mu, clusters)
```

 (Graham et al., 2017)

Summary

- Data mining and machine learning processes take significant amounts of time and are not adaptive to changing contexts
- Interactive and Anytime algorithms put the human in the loop to improve accuracy and time with significant tradeoffs
- Interruptible algorithms are designed to give the user the *option* to interactive with an algorithm with no penalties if he or she chooses not to do so.
- HOLA is a system designed to make use of an operating system architecture to manage the interrupts and visualization

Future Work

- What would a visualization system that adapts to various model and data states look like?
  - Buffering state for concurrent modification
- Incorporate git-like strategies for data experimentation
  - Branching data to allow for concurrent experiments that are independent
  - Merges and commits to persist successful changes’
- User studies
  - Analyst may not have prior knowledge of their data
  - Ensure visualization can communicate data and model state

The University of Oklahoma

 (Graham et al., 2017)

References

- Amershi, S., Fogarty, J., & Weld, D. (2012, May). Regroup: Interactive machine learning for on-demand group creation in social networks. In Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (pp. 21-30). ACM.
- Awasthi, P., Balcan, M. F., & Voevodski, K. (2014, January). Local algorithms for interactive clustering. In ICML (Vol. 14, pp. 550-558).
- Lad, S., & Parikh, D. (2014, September). Interactively guiding semi-supervised clustering via attribute-based explanations. In European Conference on Computer Vision (pp. 333-349). Springer, Cham.

References cont.

- Vlachos, M., Lin, J., Keogh, E., & Gunopulos, D. (2003). A wavelet-based anytime algorithm for k-means clustering of time series. In In Proc. Workshop on Clustering High Dimensionality Data and Its Applications.
- Ueno, K., Xi, X., Keogh, E., & Lee, D. J. (2006, December). Anytime classification using the nearest neighbor algorithm with applications to stream mining. In Data Mining, 2006. ICDM'06. Sixth International Conference on (pp. 623-632). IEEE.