Online Learning and Optimization

Introduction

Course Details

- **Name**: CS773(A) – Online Learning and Optimization
- **Nickname**: OLO
- **Instructor**: Purushottam “Puru” Kar (`purushot`)
- **Teaching Assistant**: Vijay Keswani (`vijaykes`)
- **Lectures**: TuTh 1700-1830 hrs, KD102
- **Office hours:**
  Puru: Fridays 1600-1700 hrs
  Vijay: Wednesdays 1600-1700 hrs
- **Website:** http://tinyurl.com/olo15-16w
- **Internal**: http://tinyurl.com/olo15-16wi

Auditors

- Please send a mail to Vijay confirming your decision
- Do this even if you have spoken to Vijay/Puru
- Feel free to participate in all aspects of the course
  - Attend lectures
  - Assist creditors in scribing notes
  - Submit assignments – will be graded*
  - Appear for examinations – will be graded*
  - Participate in project groups

Grading Scheme

- 15%: Assignments
  - Paper-pen (although LaTeX-ed preferred)
  - Programming-based
- 15%: Scribing lecture notes for **one** lecture
  - Typeset in LaTeX
- 15%: Mid-semester examination
- 15%: End-semester examination
- 40%: Term Project

Obtaining significant and publishable results in the project would merit an A grade irrespective of performance in other components of the course.

Scribing Duties

- Schedule up on internal website
- Can swap lectures with others
  - Please inform Vijay and Puru beforehand
- Use the prescribed style file
  - Available on internal website
  - Do not edit style file – ask Puru in case of doubts
- Take pride in your scribed notes
  - Well explained, details worked out
  - Well referenced, proper citations, acknowledgements
  - Properly formatted definitions, theorems, lemmata
  - Illustrations when necessary
  - Sample scribe also present

Project

- Form groups of 2 (1 or 3 allowed as special case)
  - Auditors can join project groups but wont be counted
  - The class can express any concerns regarding this rule
  - Make groups known to Vijay and Puru
- Project proposals (written) due before class 19th Jan
- Mid-term presentations: 1st Mar, 2016
- Final presentations + report: 12th + 14th Apr, 2016
- Breakup
  - Project proposal: 5%
  - Mid-term presentation: 10%
  - Final presentation: 20%
  - Report: 5%

Project

- Some project ideas to be put up on internal website
  - Expect list of suggested idea before 2nd Jan
  - Discuss with friends, Puru for more ideas
  - Please do not wait till 19th Jan to discuss
- Project topic needs to be related to the course
- Project has to be substantial
  - Simple implementation of existing algorithms wont do
  - Reading projects possible but require extensive coverage and insight into what was done and what *can* be done
- Objective of the course
  - Lectures act as enablers – introducing basics, tools
  - Project investigation is where thorough instruction takes place

Reference Material

- No textbook for the course
- Reference list up on website
- Locally cached copies for some on internal website

- [**BVB**] Boyd and Vandenberghe. *Convex Optimization*.
- [**BCB**] Bubeck and Cesa-Bianchi. *Regret Analysis of Stochastic and Nonstochastic Multi-armed Bandit Problems*.
- [**CBL**] Cesa-Bianchi and Lugosi. *Prediction, Learning, and Games*.
- [**HZN**] Hazan. *Introduction to Online Convex Optimization*.
- [**MRT**] Mohri, Rostamizadeh, and Talwalkar. *Foundations of Machine Learning*.
- [**SSS**] Shalev-Shwartz. *Online Learning and Online Convex Optimization*.

Use of Unfair Means

- The following are prohibited – severe penalties
  - Copying answers in pen-paper assignments
  - Copying code in programming assignments
  - Passing off known results as one's own
  - Manipulating experimental results

- The following are prohibited – credit deductions
  - Using material in scribes (figures, text) without acknowledging
  - Using help from auditors in projects without acknowledging

Machine Learning 101

Machine Learning

Machine Learning

“The art and science of designing adaptive algorithms”

Machine Learning

"The art and science of designing adaptive algorithms"

Spam Filtering

A different "classification problem" for every individual

A different problem for every context

Machine Learning

"The art and science of designing adaptive algorithms"

Spam Filtering

A different "classification problem" for every individual

A different problem for every context

Machine Learning

“The art and science of designing adaptive algorithms”

Spam Filtering

A different “classification problem” for every individual

A different problem for every context

Subject:	[all] New Pizza Counter at New SAC
From:	"DOSA" <dosa@iitk.ac.in>
Date:	Wed, October 28, 2015 10:07 am
To:	all@lists.iitk.ac.in
Cc:	dosa@iitk.ac.in (more)
Priority:	Normal
Options:	View Full Header | View Printable Version | Download this as a file

Subject:	[all] Lost and Found
From:	"DOSA" <dosa@iitk.ac.in>
Date:	Wed, October 28, 2015 10:07 am
To:	all@lists.iitk.ac.in
Cc:	dosa@iitk.ac.in (more)
Priority:	Normal
Options:	View Full Header | View Printable Version | Download this as a file

Machine Learning

“The art and science of designing adaptive algorithms”

Self-driving Cars

A different control problem for every locale

www.youtube.com, www.ci.pasadena.ca.us, www.newindianexpress.com, inextlive.jagran.com

Machine Learning

"The art and science of designing adaptive algorithms"

Self-driving Cars

A different control problem for every locale

www.youtube.com, www.ci.pasadena.ca.us, www.newindianexpress.com, inextlive.jagran.com

Traditional Machine Learning Primitives

Traditional Machine Learning Primitives

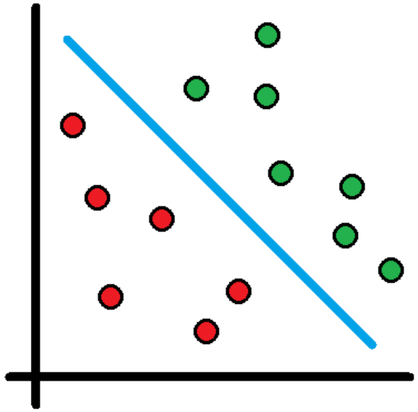

Binary Classification

Traditional Machine Learning Primitives

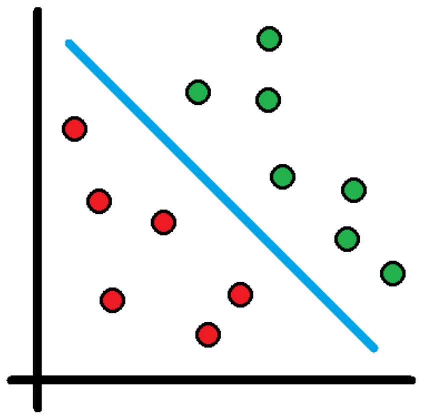

Binary Classification

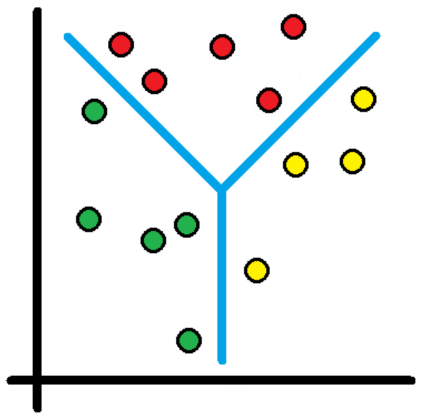

Multi Classification

Traditional Machine Learning Primitives

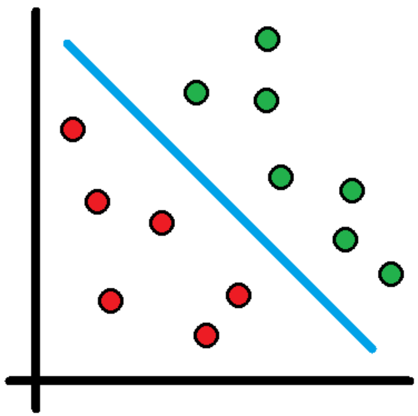

Binary Classification

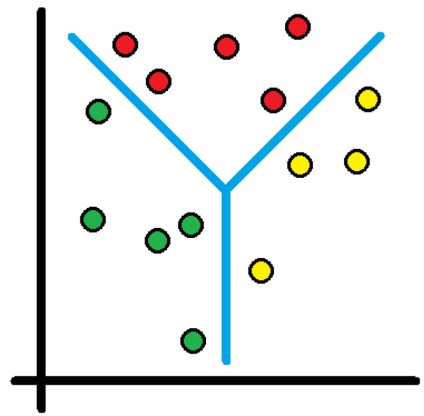

Multi Classification

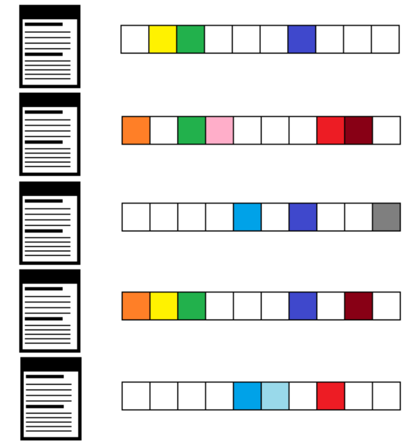

Multi-label Classfn

Traditional Machine Learning Primitives

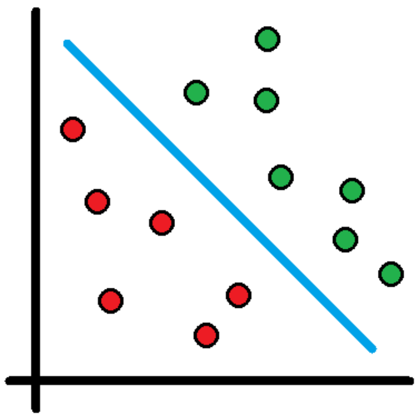

Binary Classification

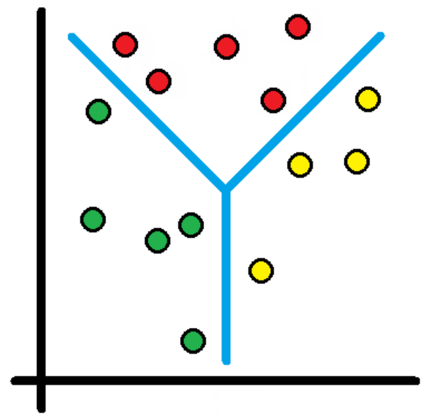

Multi Classification

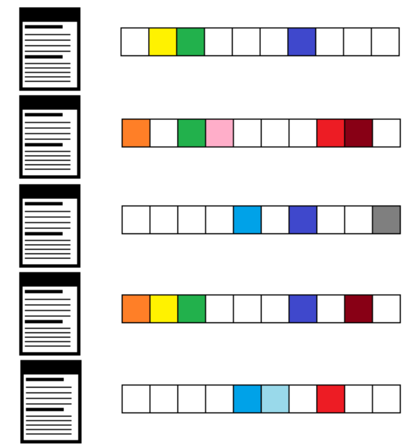

Multi-label Classfn

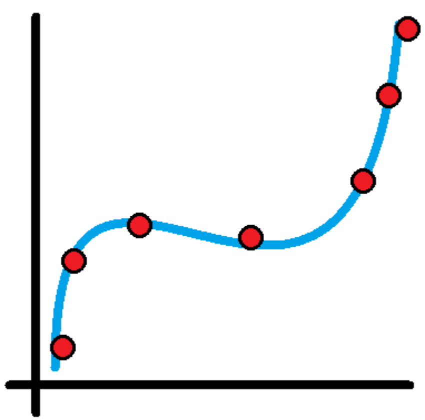

Regression

Traditional Machine Learning Primitives

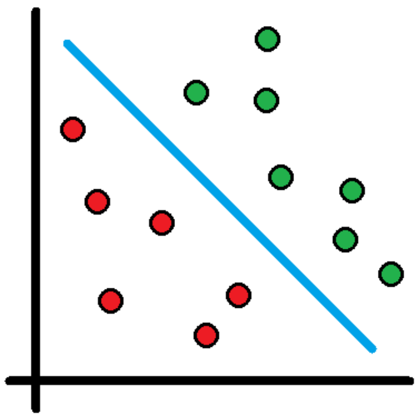

Binary Classification

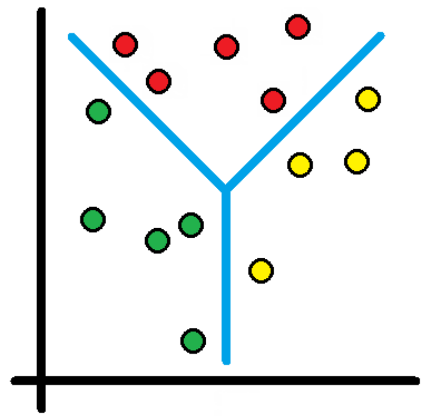

Multi Classification

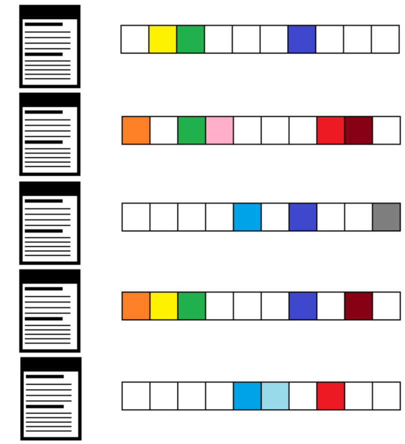

Multi-label Classfn

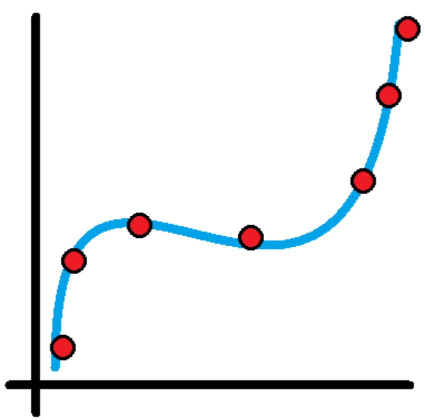

Regression

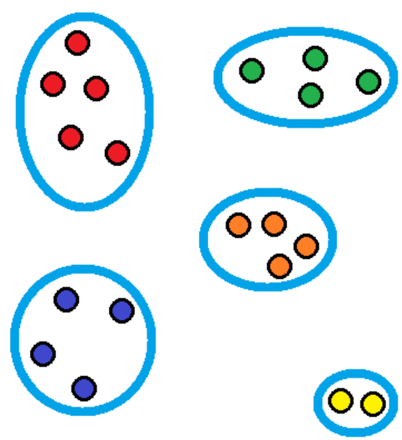

Clustering

Traditional Machine Learning Primitives

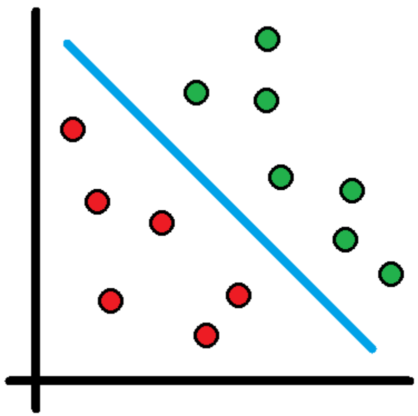

Binary Classification

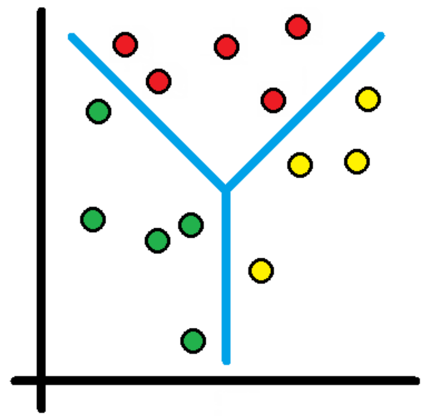

Multi Classification

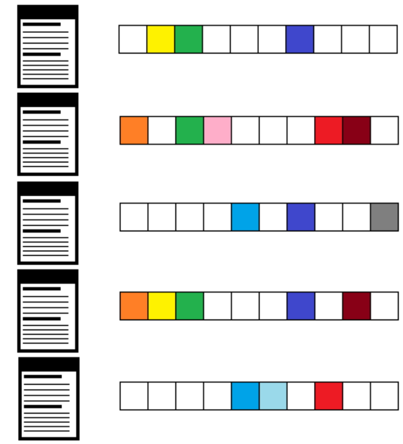

Multi-label Classfn

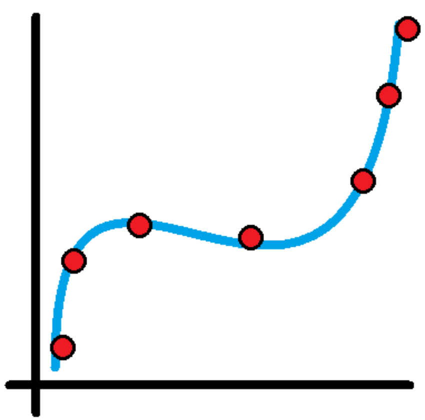

Regression

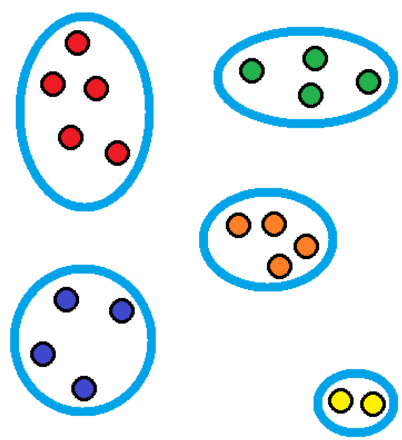

Clustering

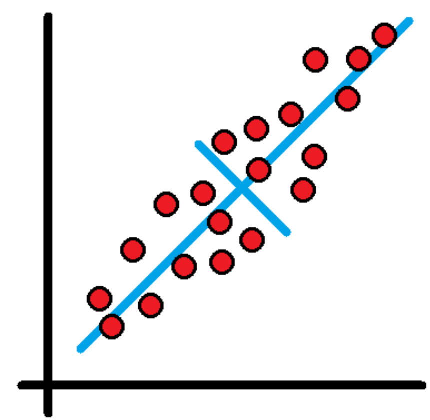

Component Analysis

Machine Learning – Perspectives

www.iconarchive.com

Machine Learning – Perspectives

www.iconarchive.com

Machine Learning – Perspectives

www.iconarchive.com

Machine Learning – Perspectives

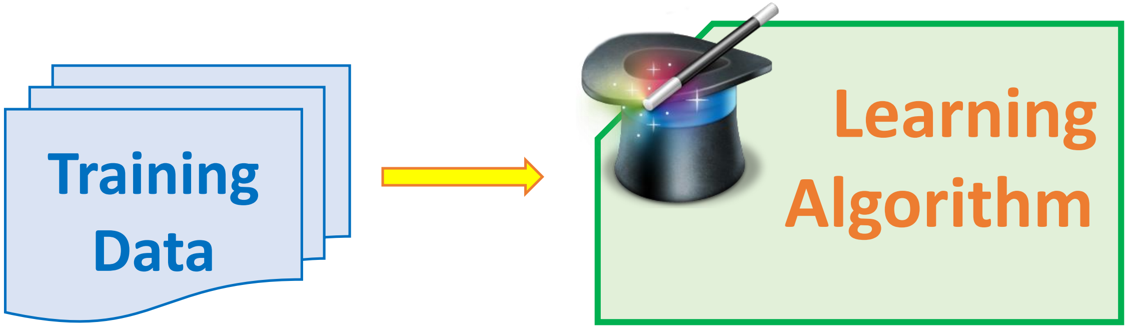

www.iconarchive.com

Machine Learning – Perspectives

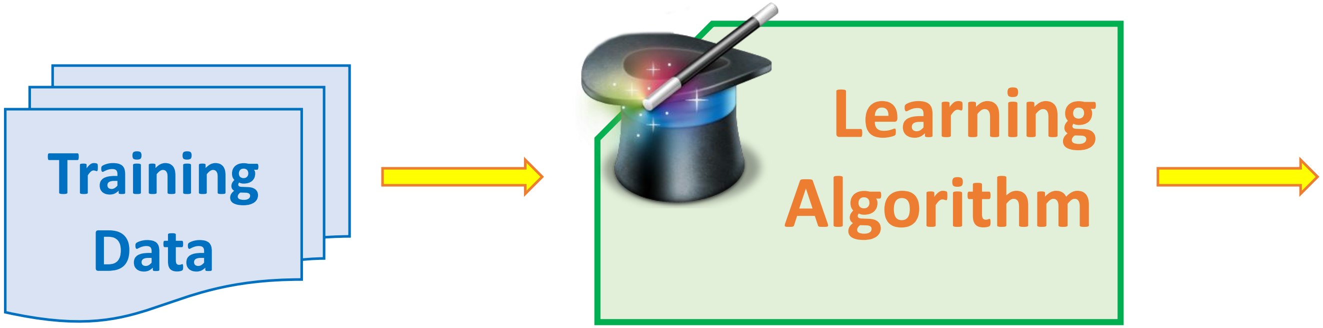

www.iconarchive.com

Machine Learning – Perspectives

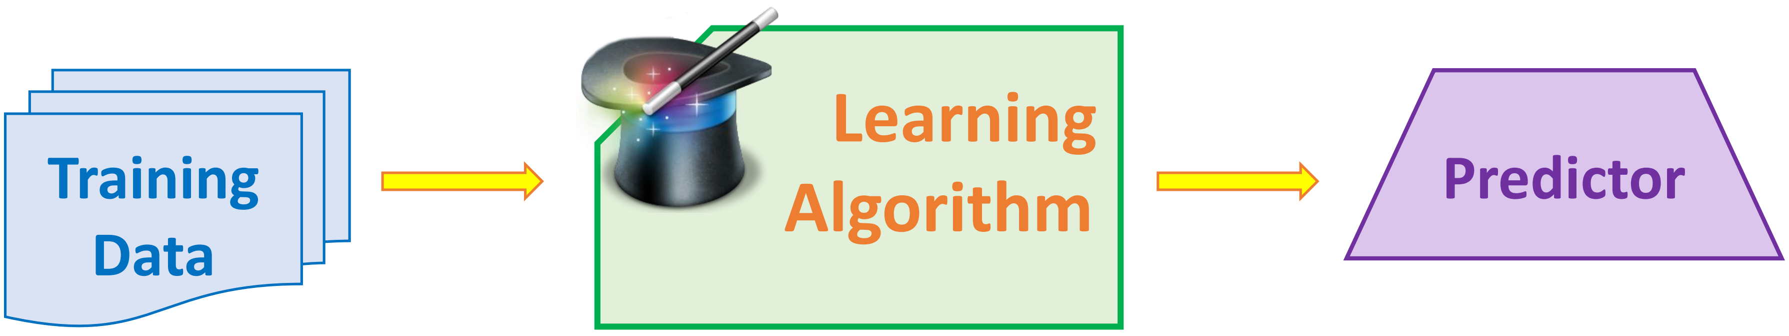

www.iconarchive.com

Machine Learning – Perspectives

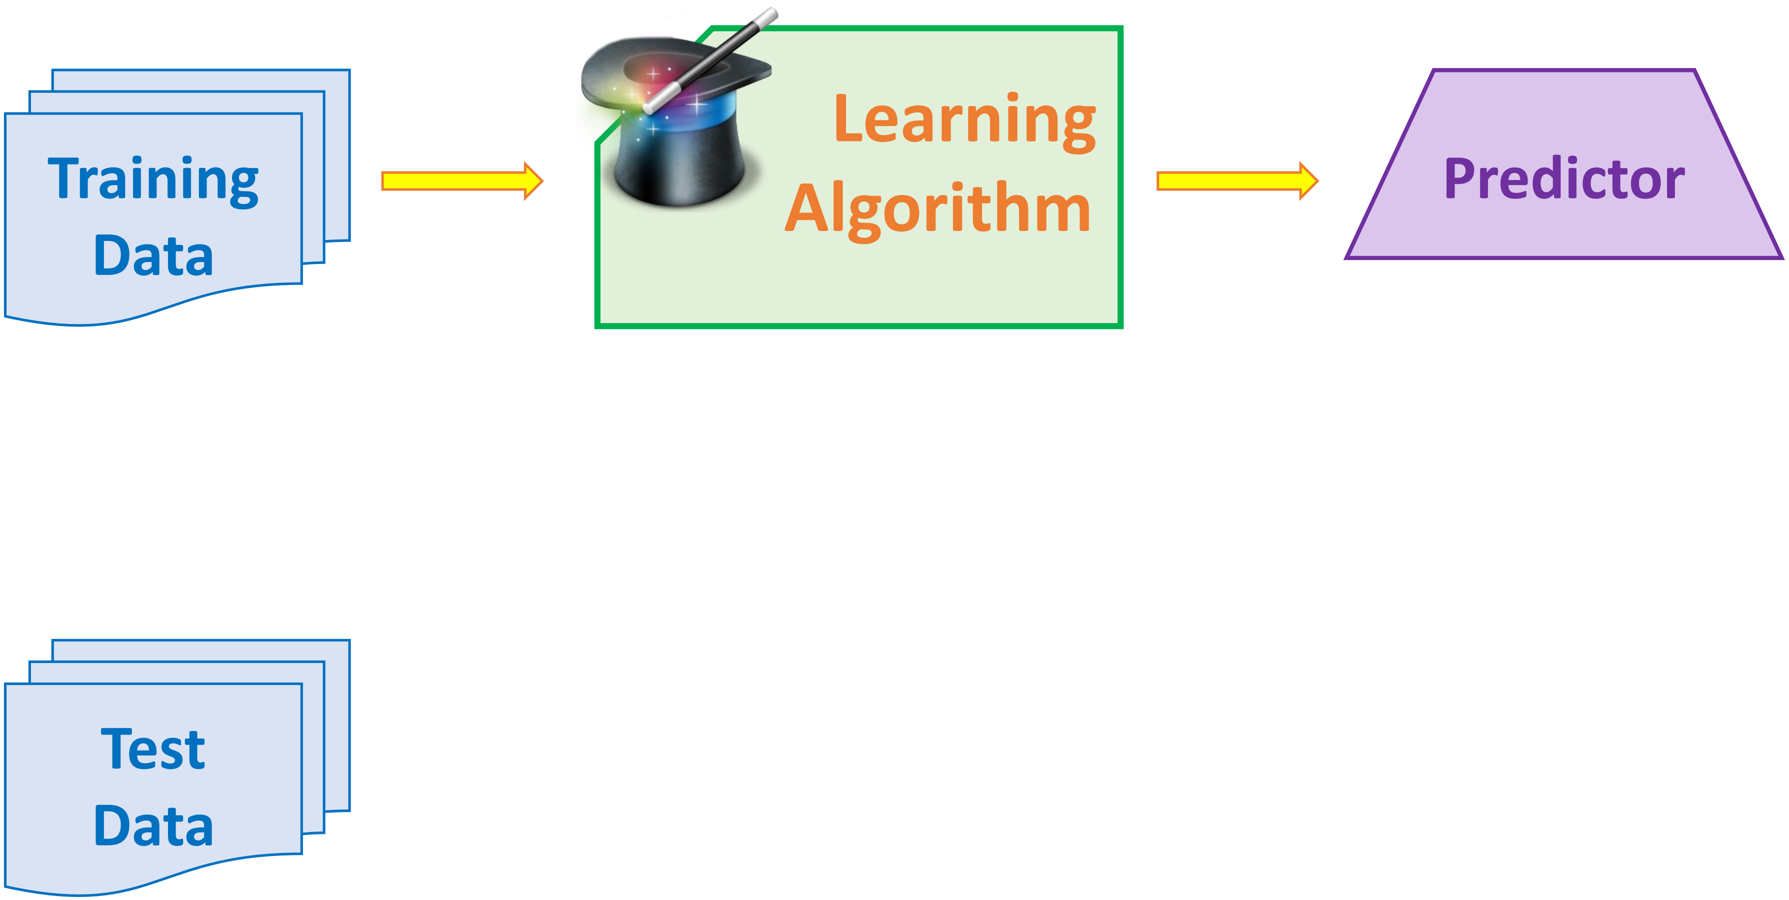

www.iconarchive.com

Machine Learning – Perspectives

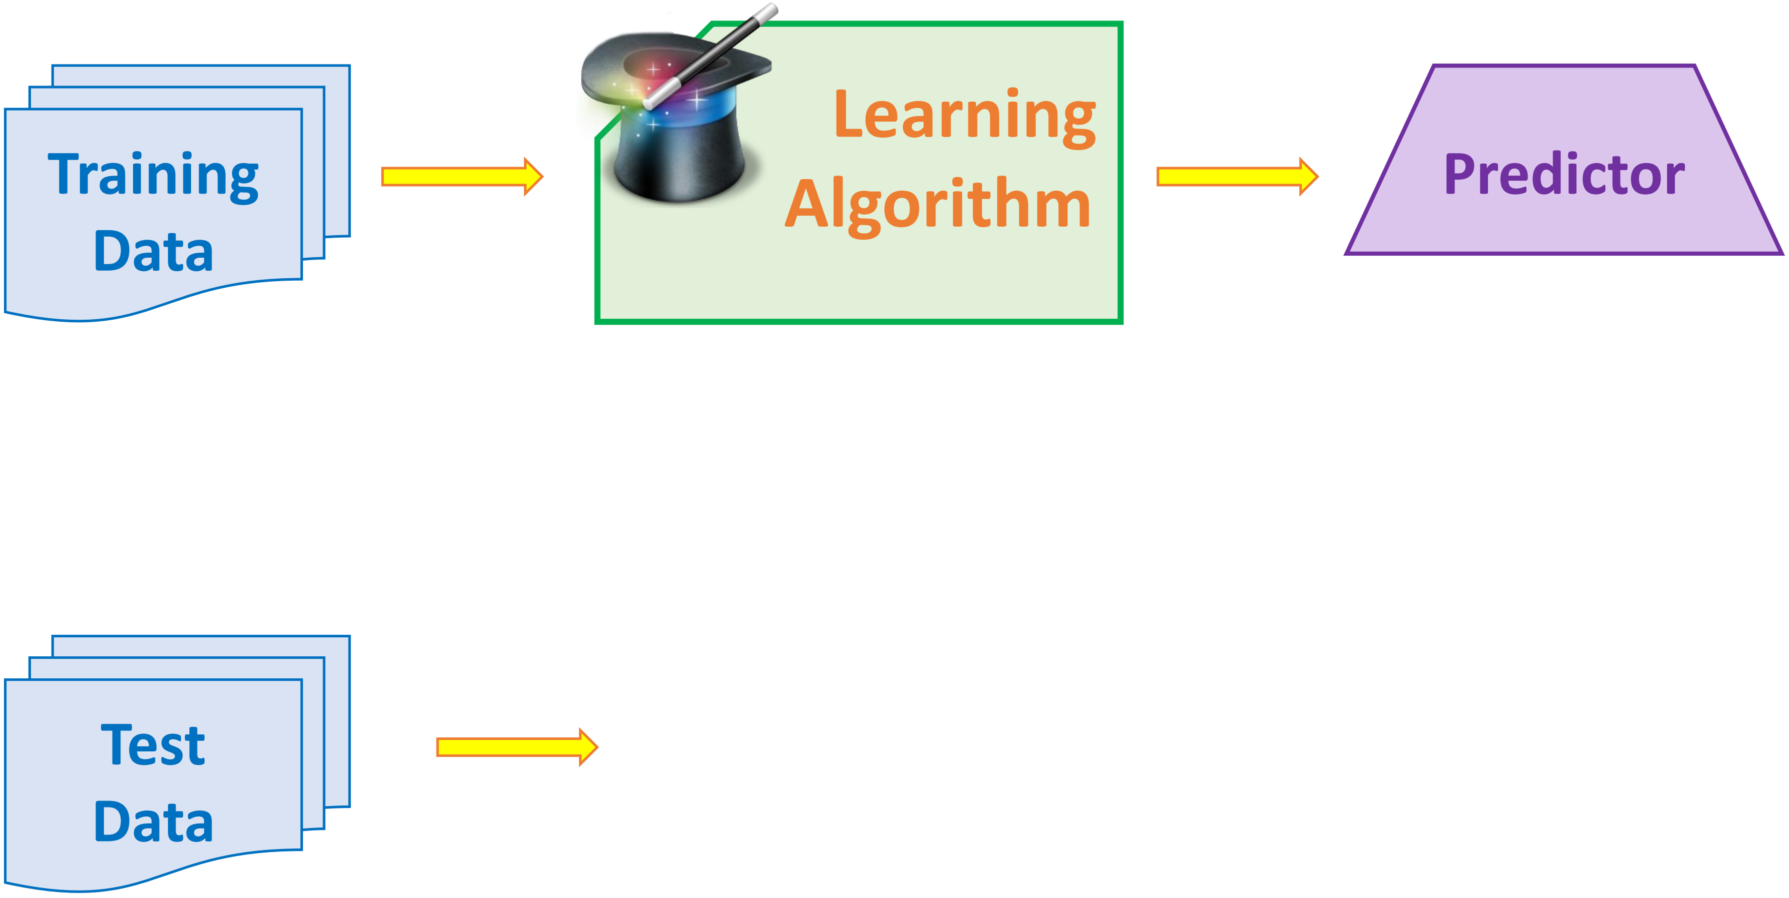

www.iconarchive.com

Machine Learning – Perspectives

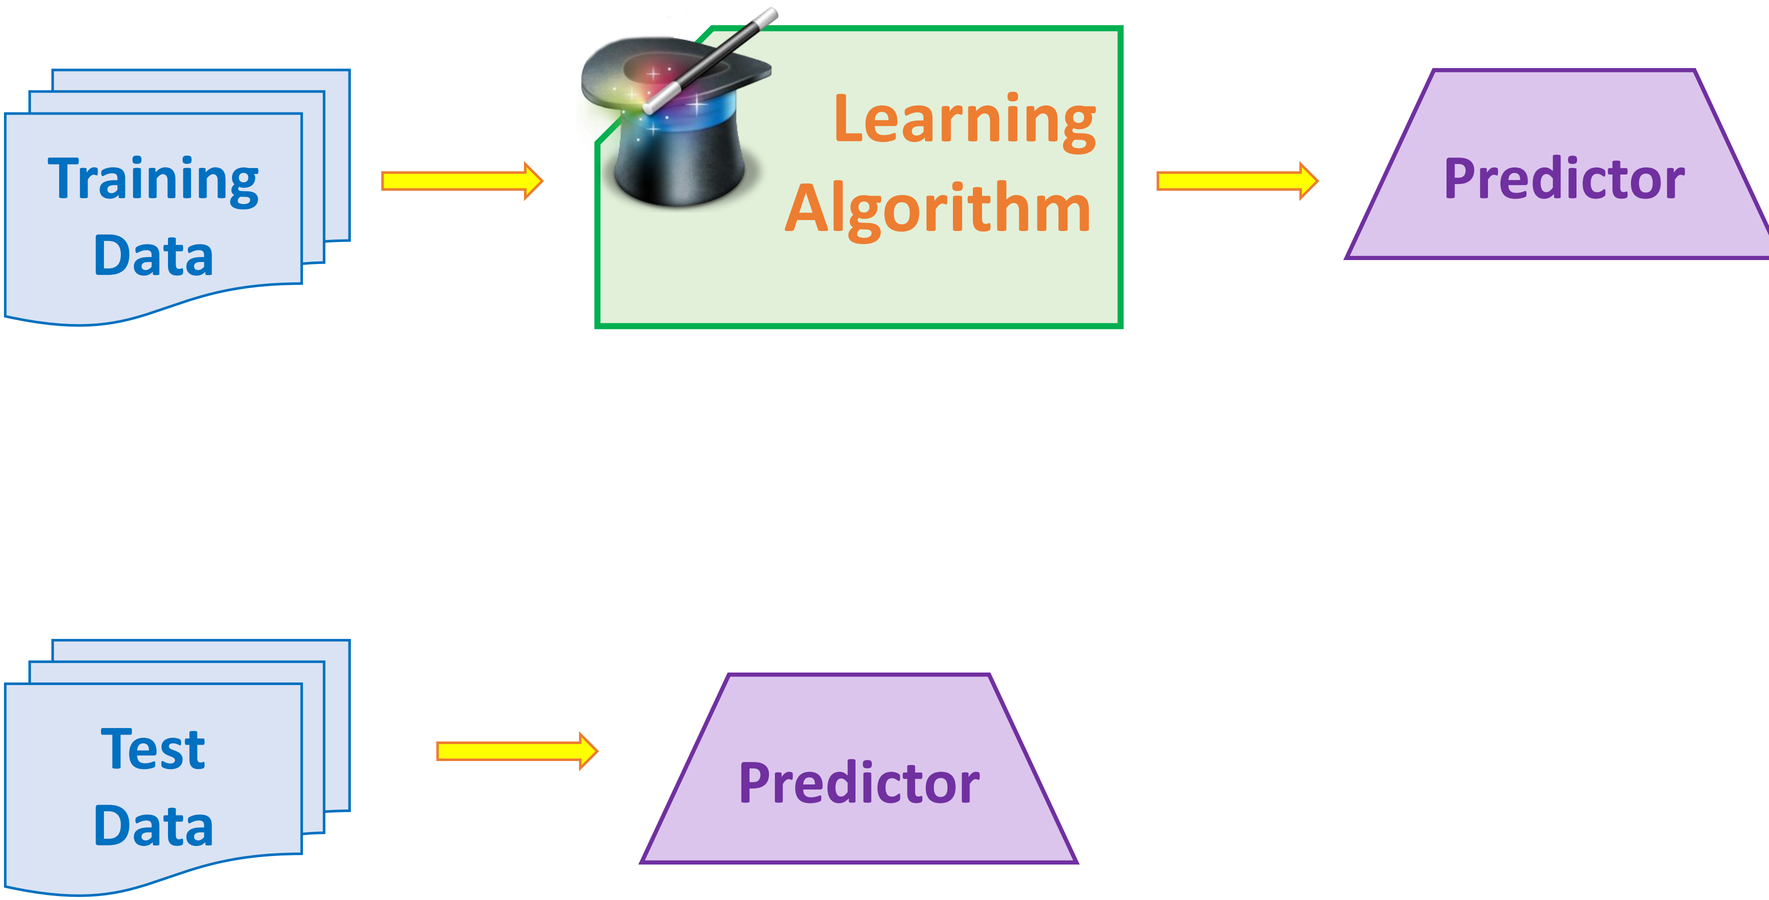

www.iconarchive.com

Machine Learning – Perspectives

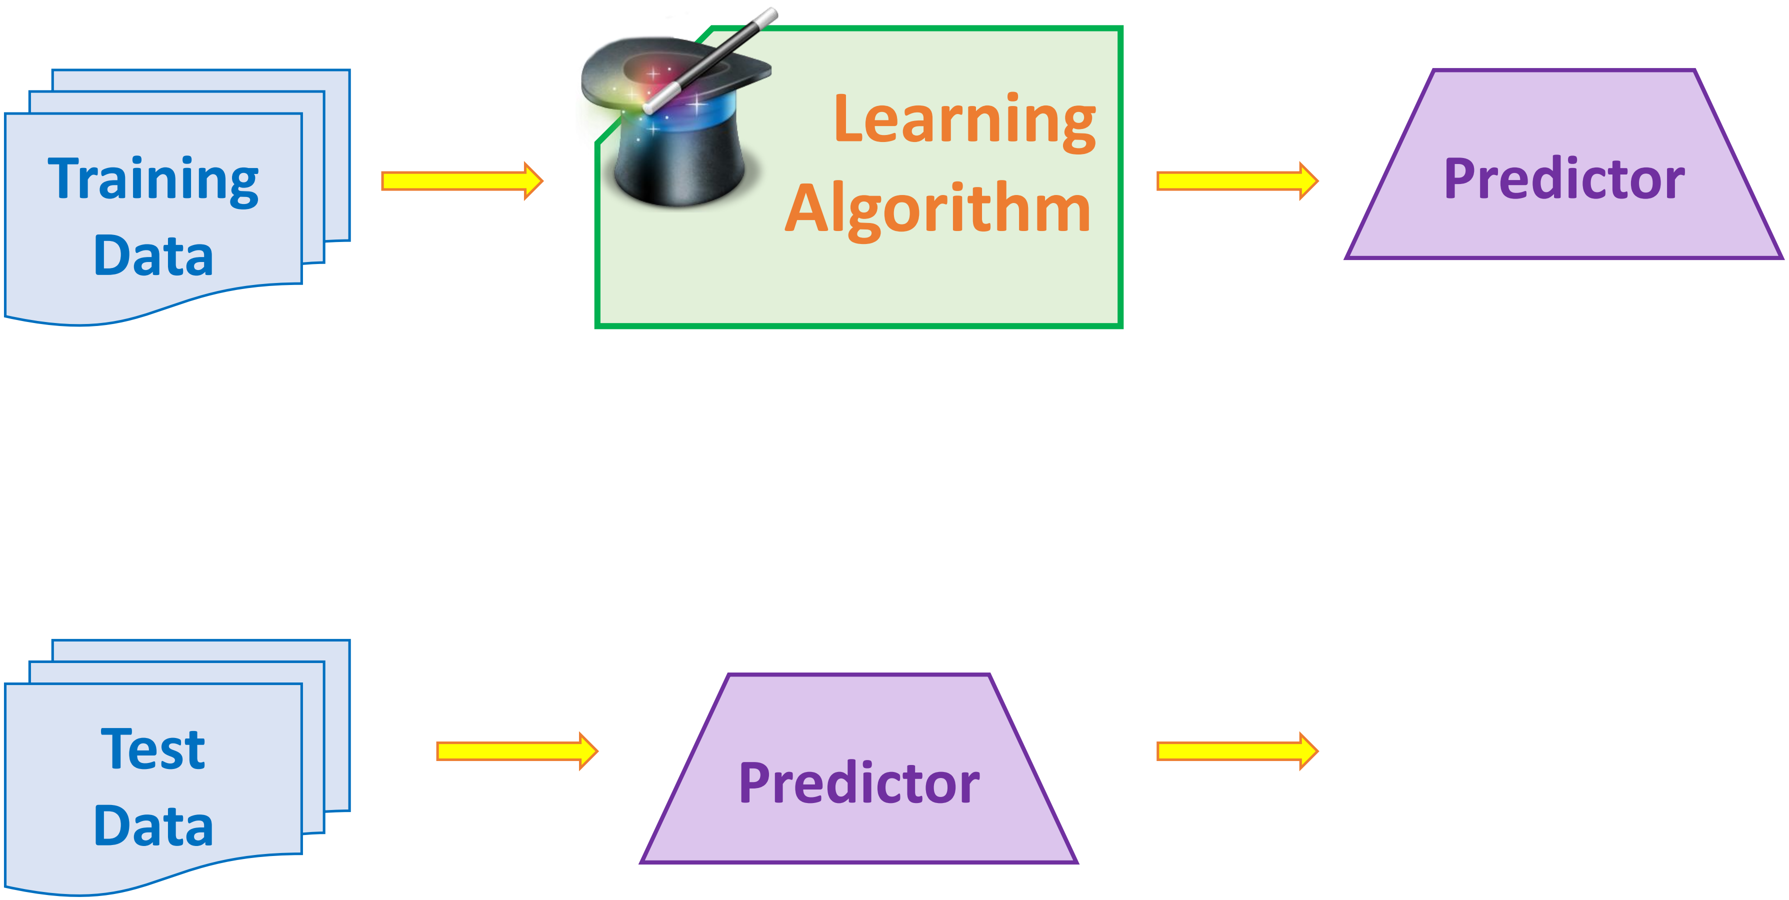

www.iconarchive.com

Machine Learning – Perspectives

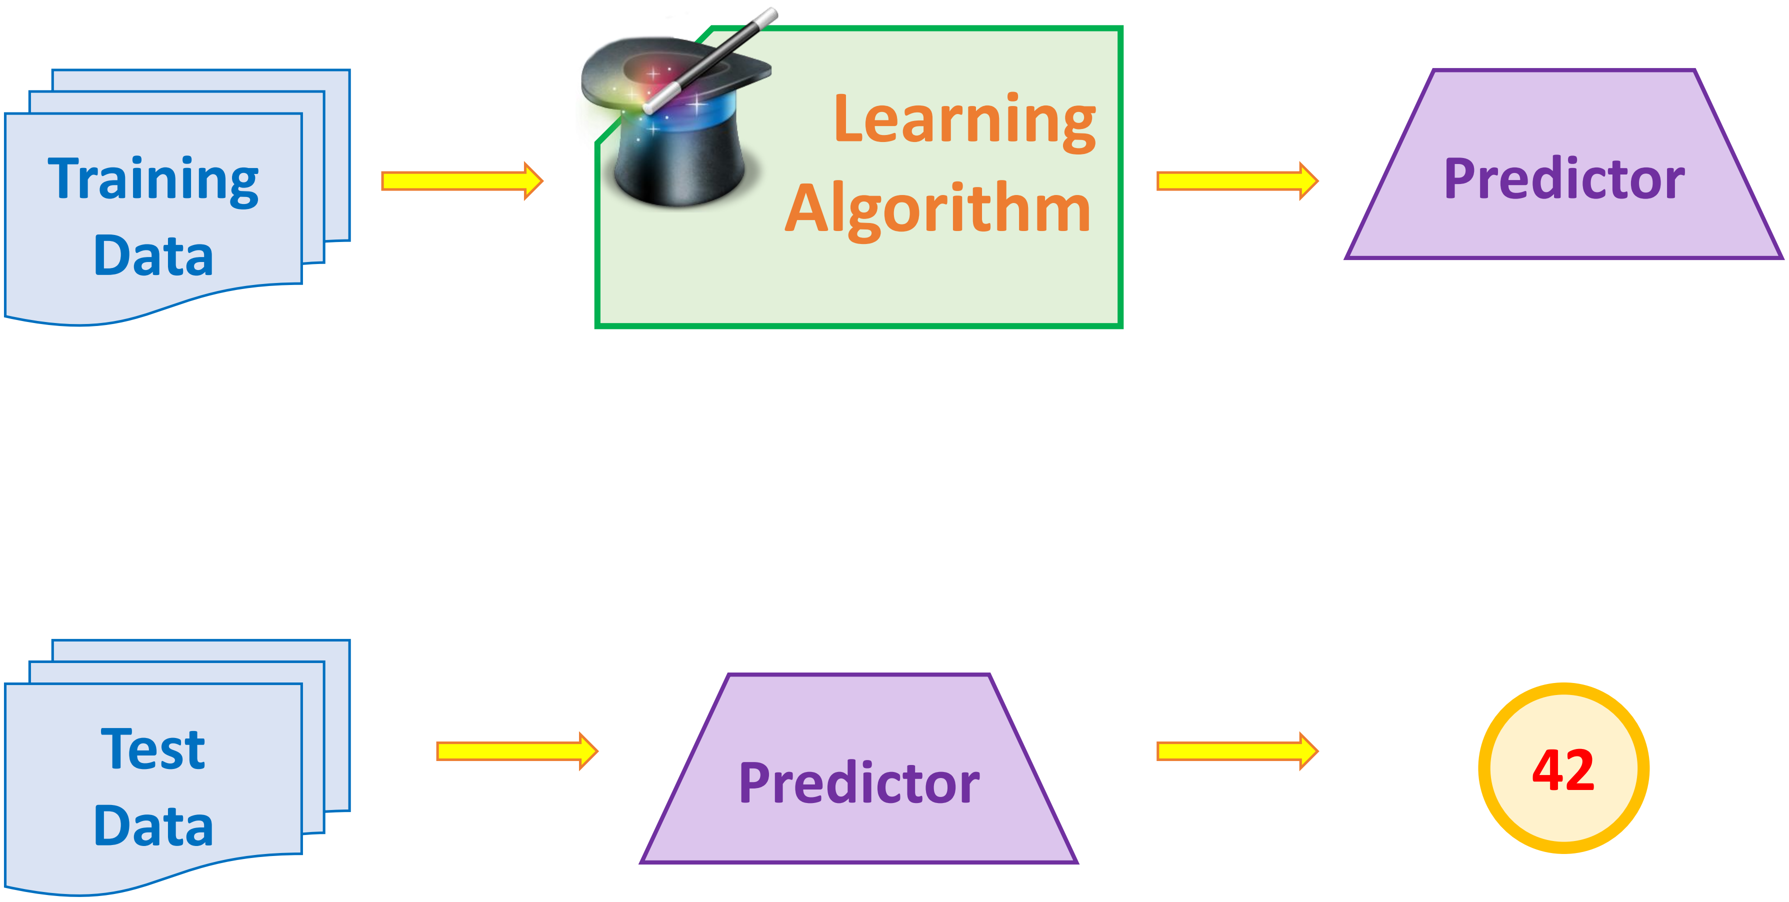

www.iconarchive.com

Machine Learning in Action

Linear Binary Classification

Machine Learning in Action

Linear Binary
Classification

Machine Learning in Action

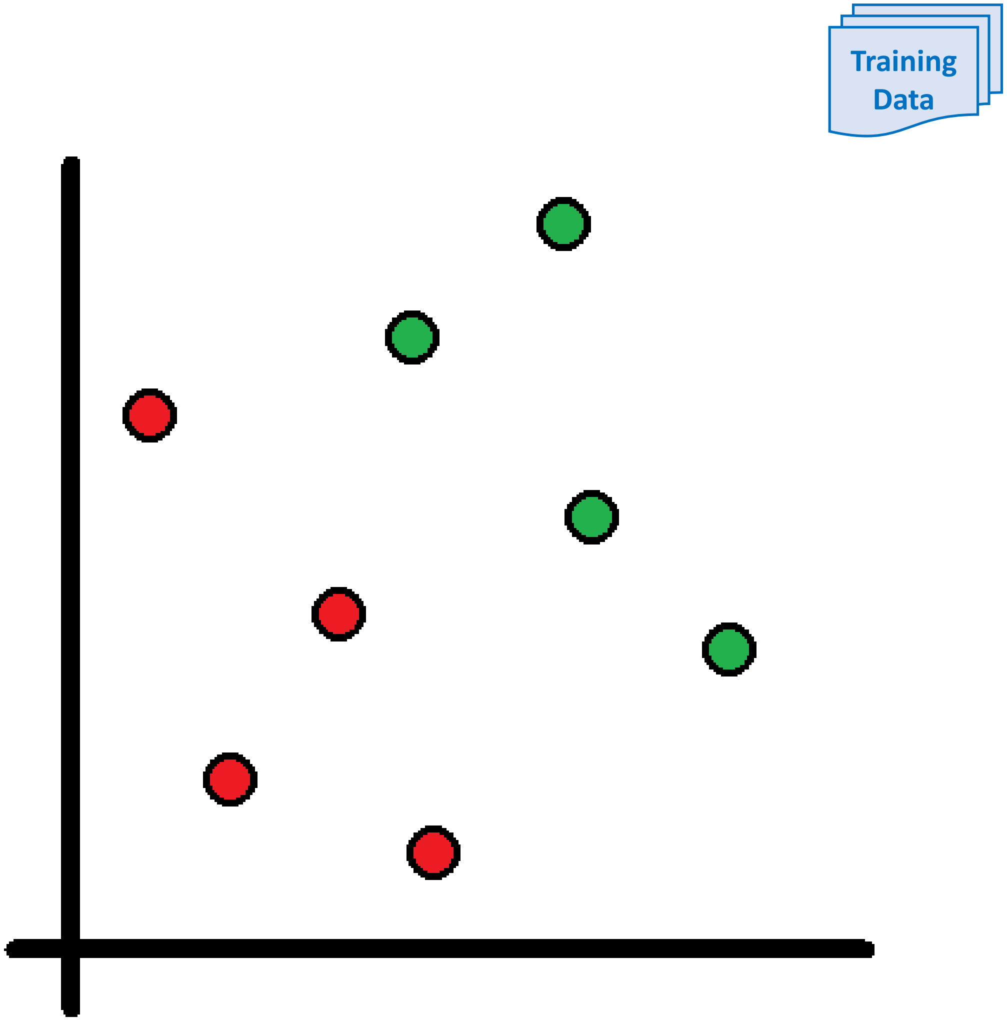

Linear Binary
Classification

Machine Learning in Action

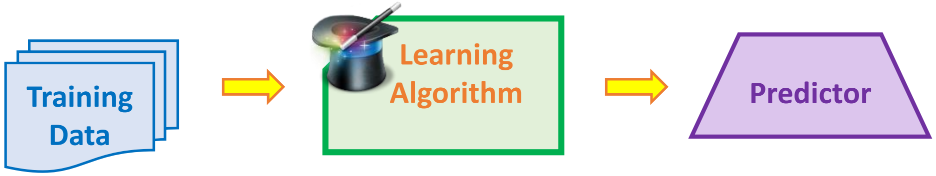

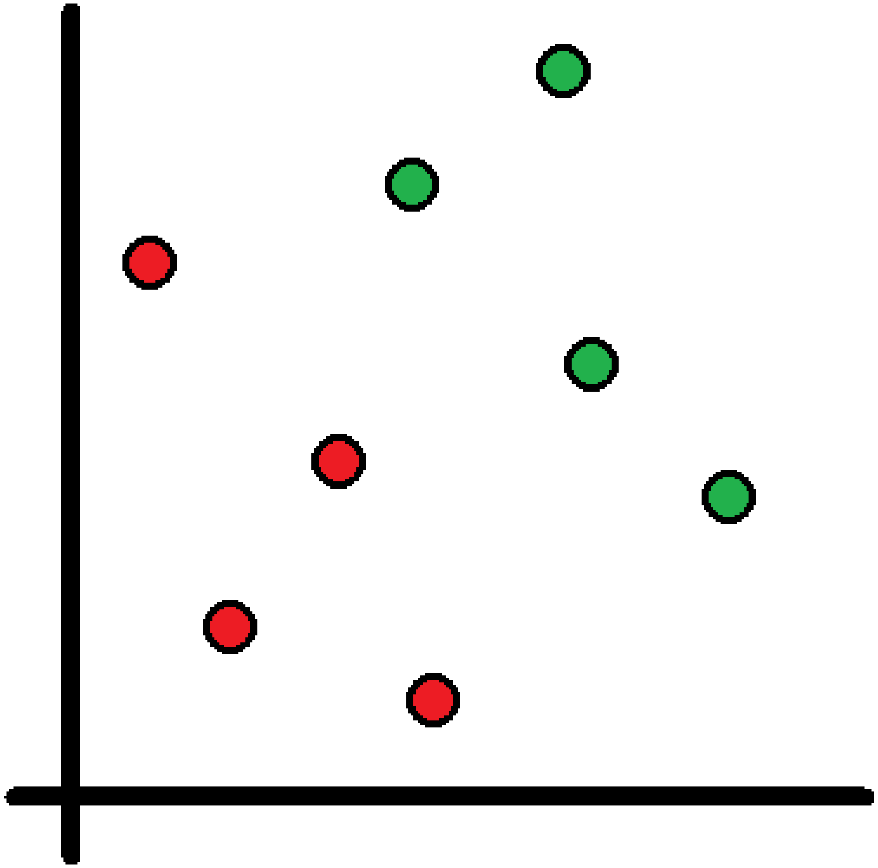

Linear Binary Classification

Machine Learning in Action

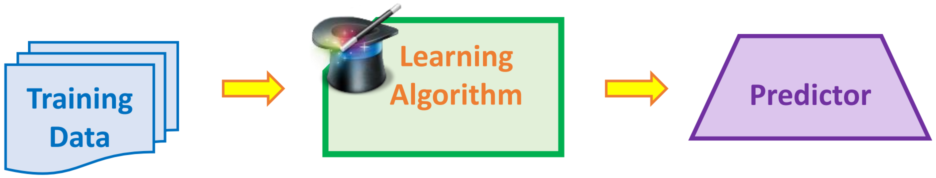

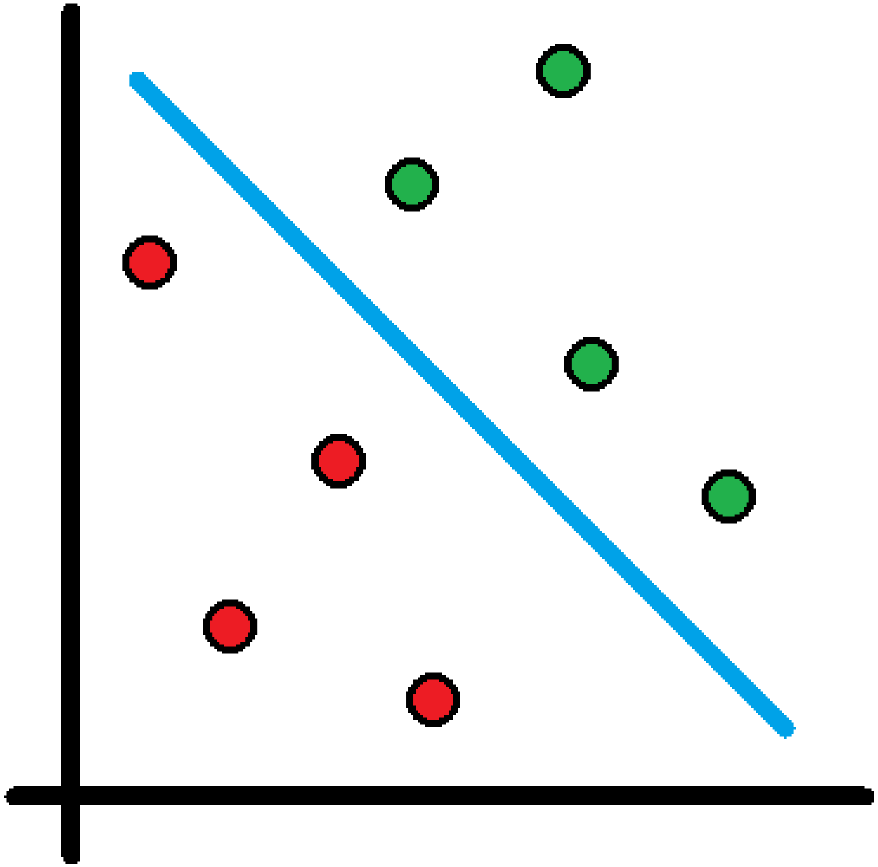

Linear Binary Classification

Machine Learning in Action

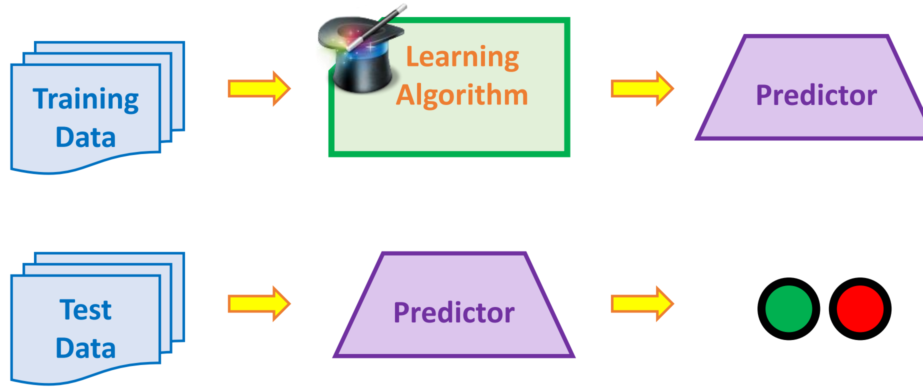

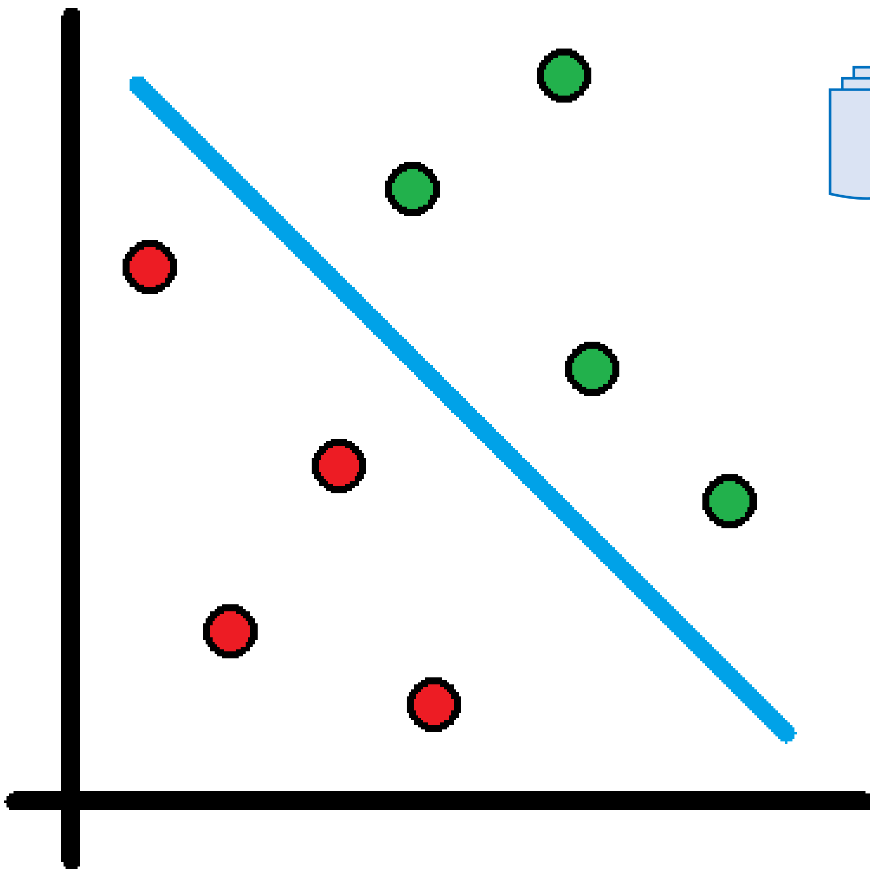

Linear Binary Classification

Machine Learning in Action

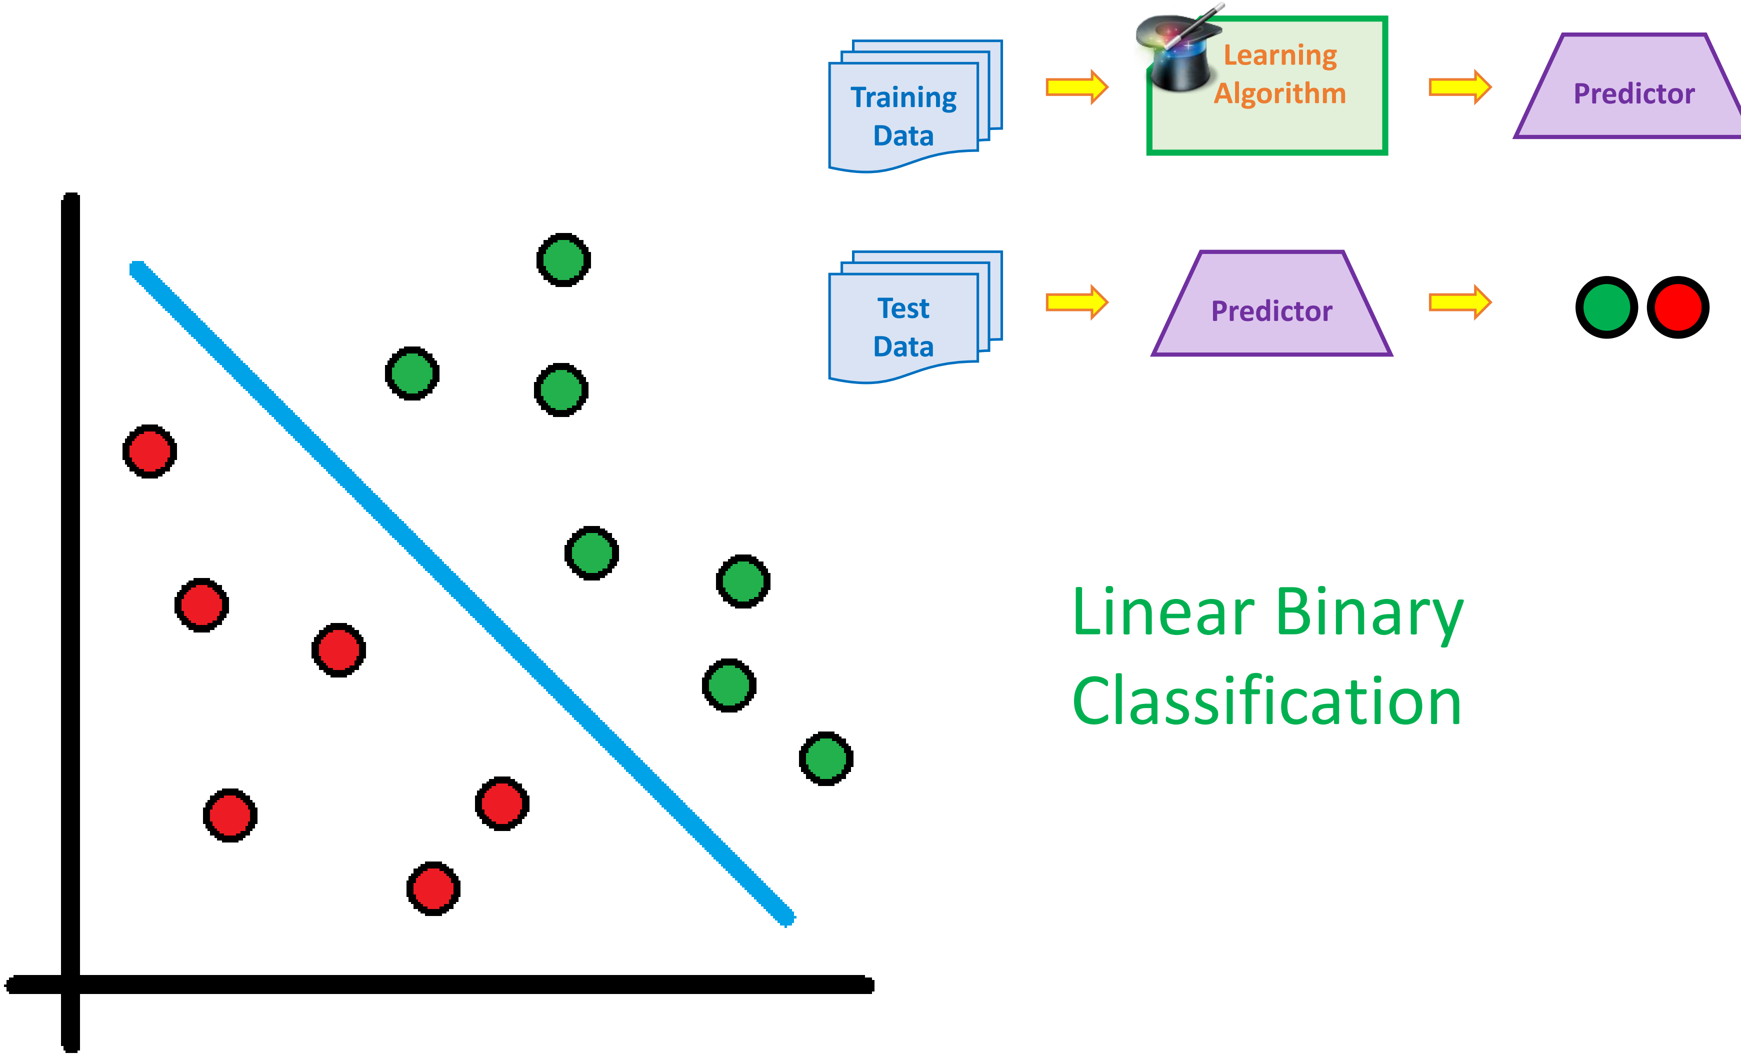

Linear Binary Classification

Machine Learning – Perspectives

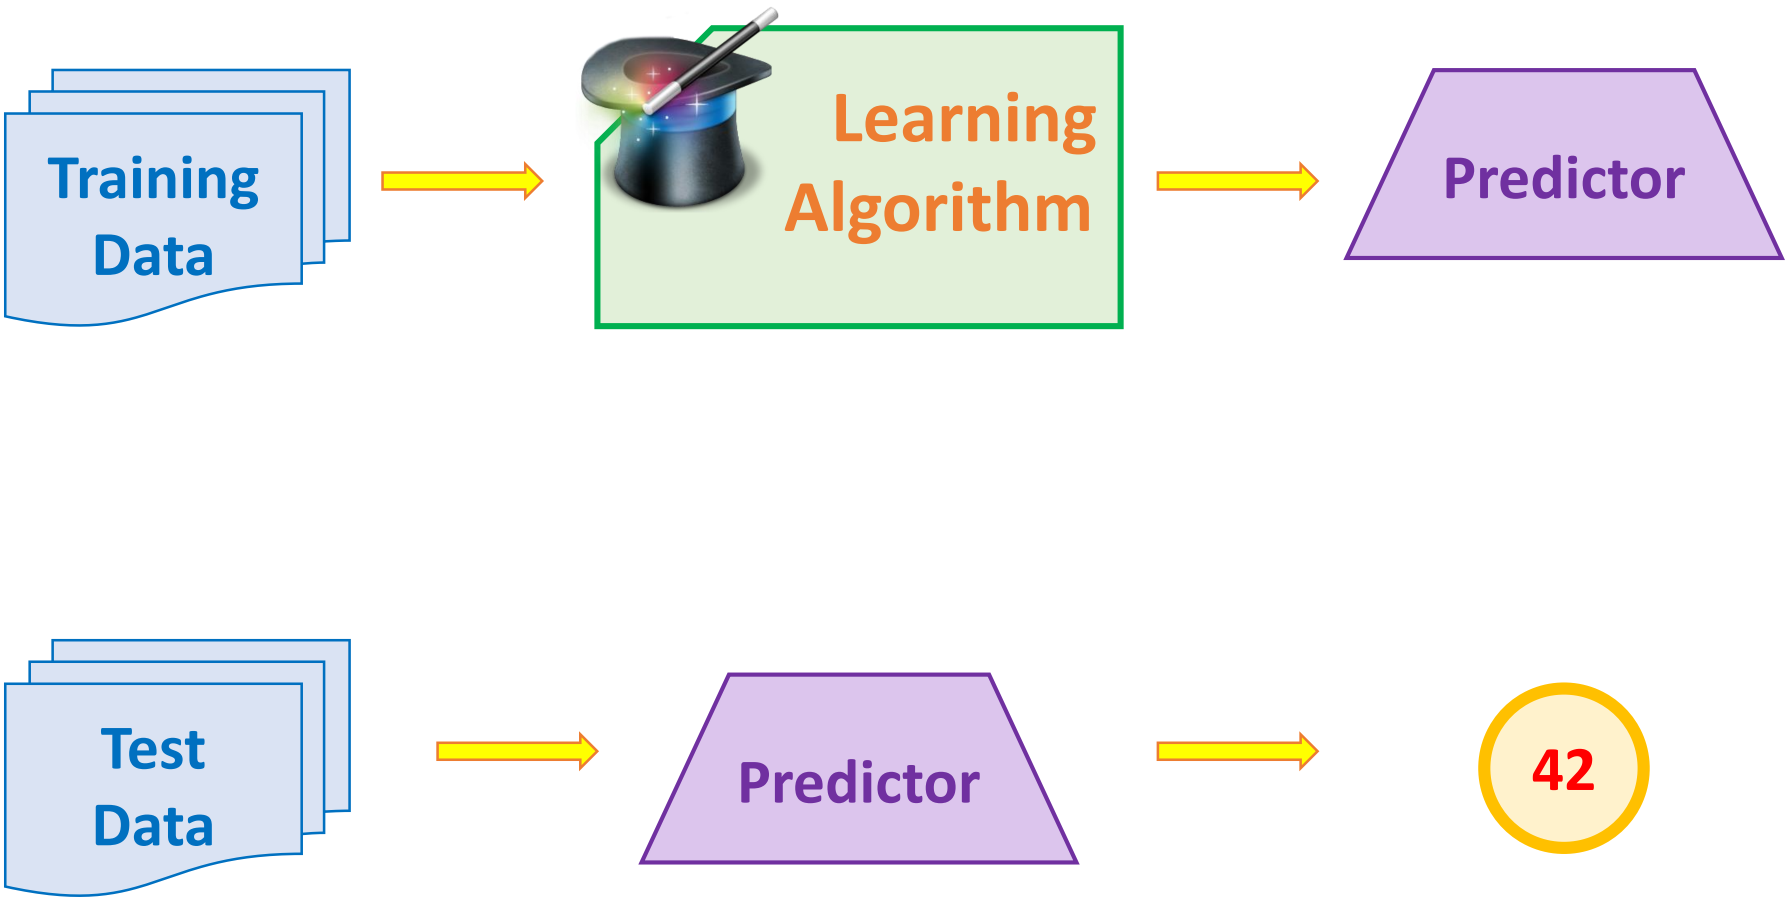

Machine Learning – Perspectives

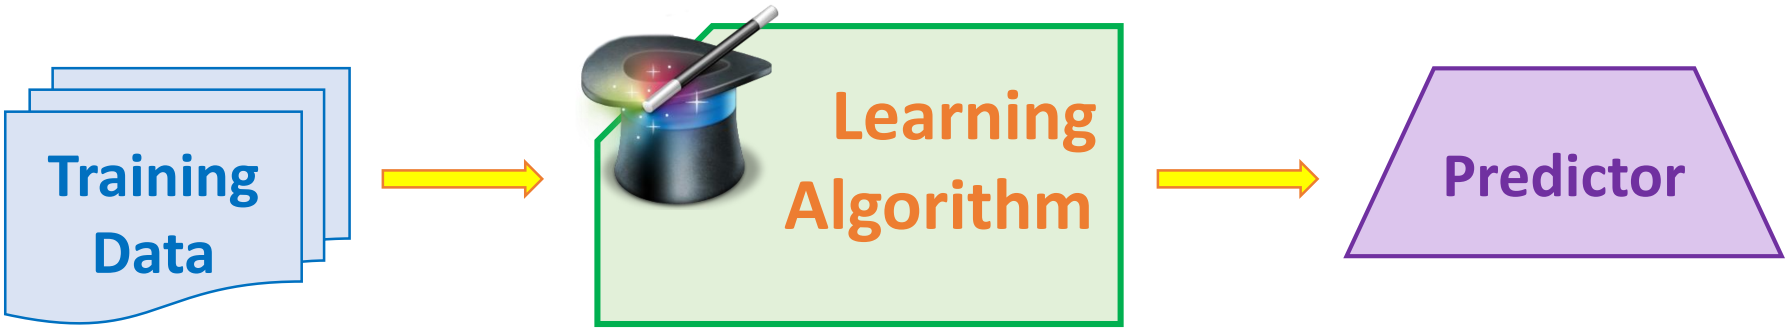

Machine Learning – Perspectives

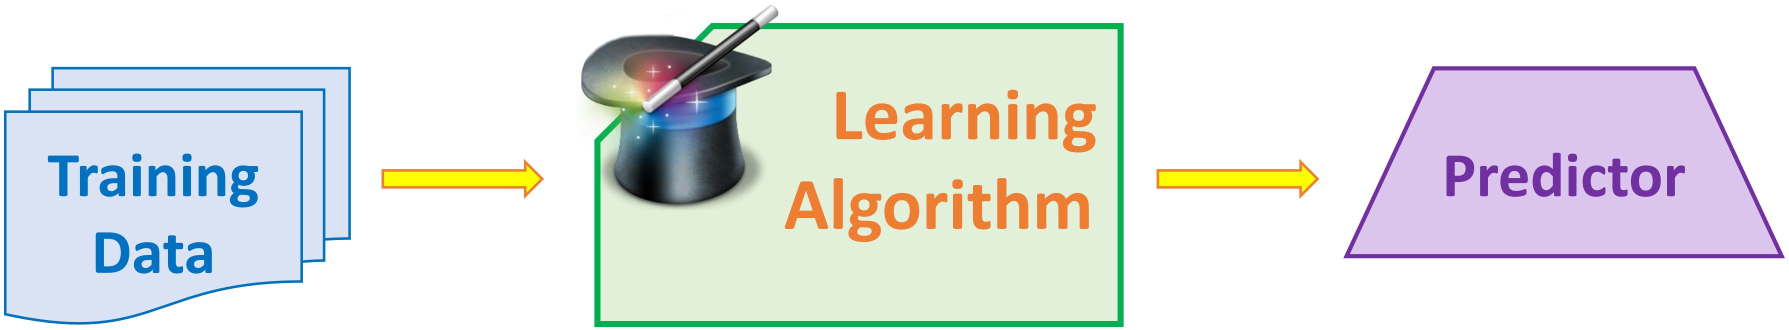

Design Questions

Machine Learning – Perspectives

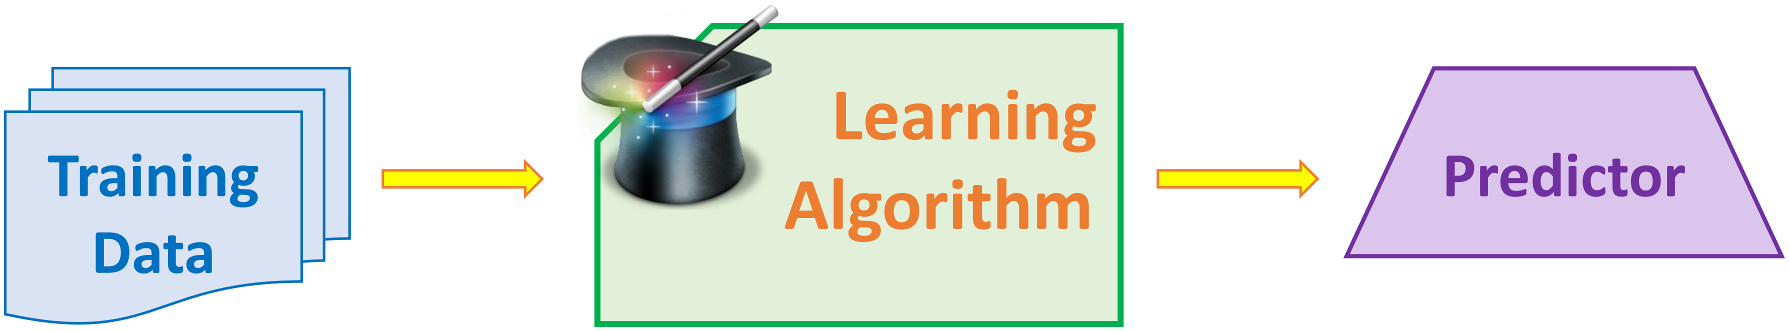

Design Questions

- How is training data presented/acquired?

Machine Learning – Perspectives

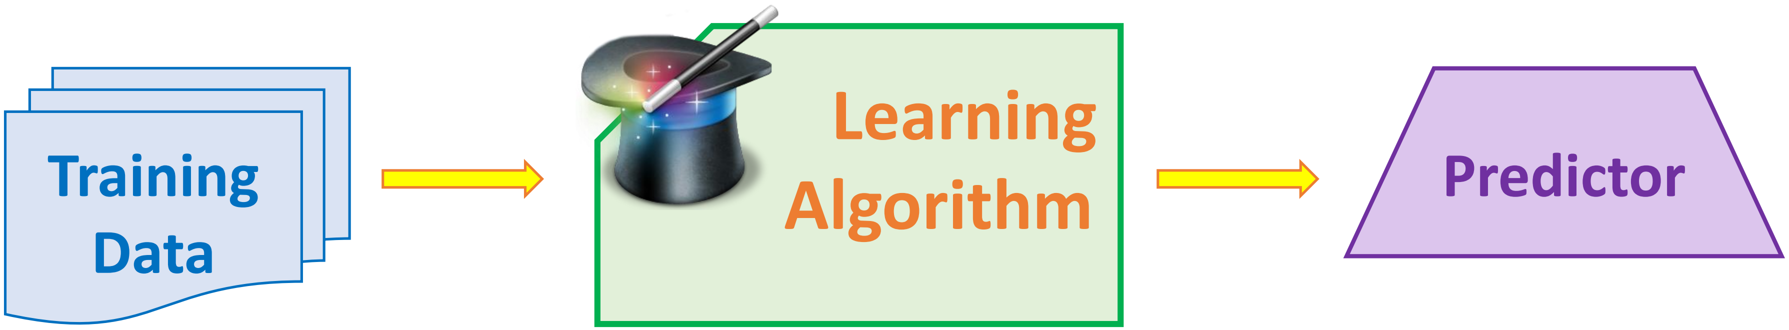

Design Questions

- How is training data presented/acquired?

- What is the predictor supposed to do?

Machine Learning – Perspectives

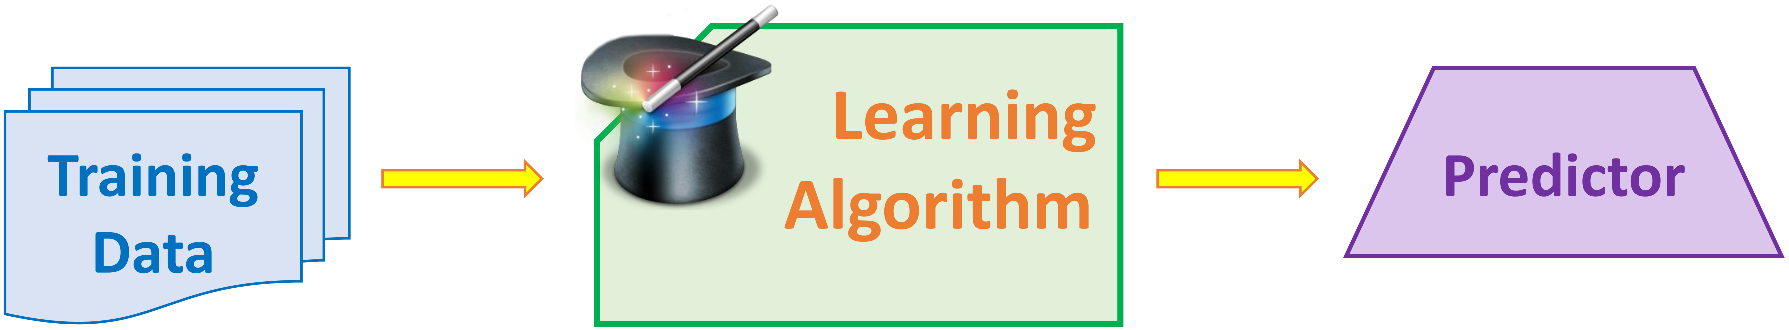

Design Questions

- How is training data presented/acquired?

- What is the predictor supposed to do?
  - Labels: binary, k-ary, multiset, real number, natural numbers
  - Assignments: new representation, categorization, permutation

Machine Learning – Perspectives

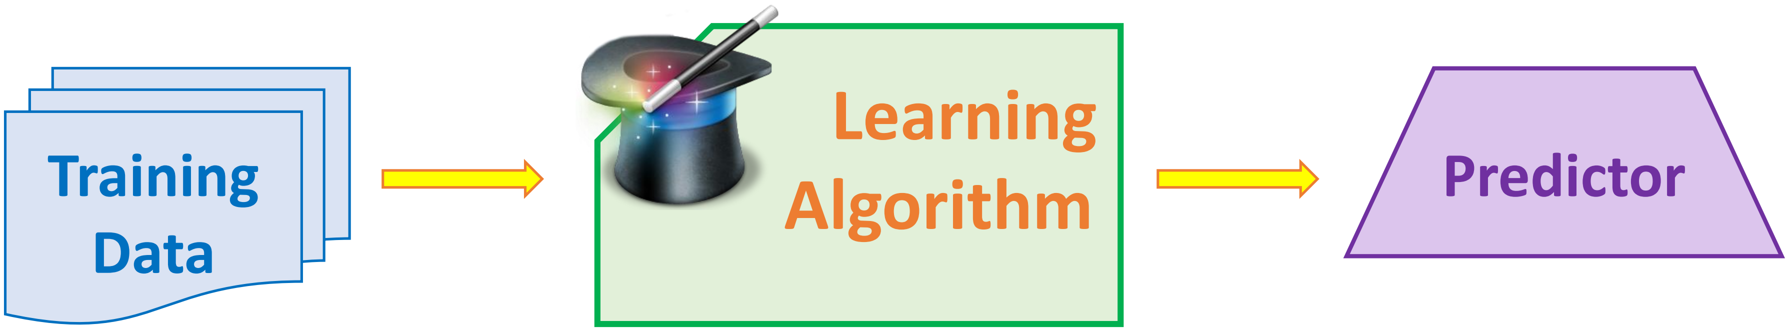

Design Questions

- How is training data presented/acquired?
  - Fully labelled, partially labelled, label-on-request?
  - All at once, one at a time, upon request?
  - Generated passively, noisily, adversarially, by an MDP?
- What is the predictor supposed to do?
  - Labels: binary, k-ary, multiset, real number, natural numbers
  - Assignments: new representation, categorization, permutation

Machine Learning – Perspectives

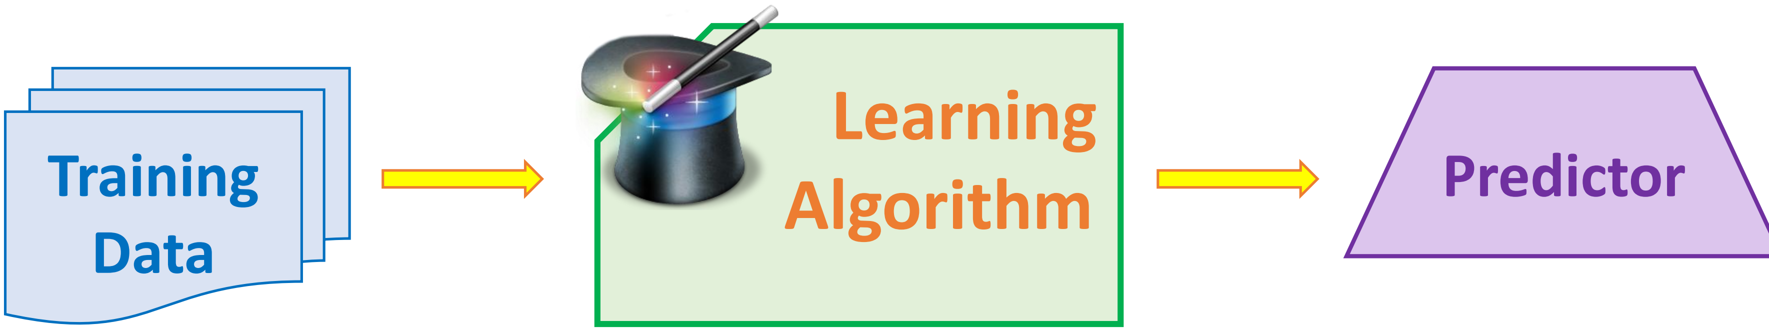

Design Questions

- How is training data presented/acquired?
  - Fully labelled, partially labelled, label-on-request?
  - All at once, one at a time, upon request?
  - Generated passively, noisily, adversarially, by an MDP?
- What is the predictor supposed to do?
  - Labels: binary, k-ary, multiset, real number, natural numbers
  - Assignments: new representation, categorization, permutation

Online Learning

Traditional “Batch” Learning

pinterest.com

Traditional "Batch" Learning

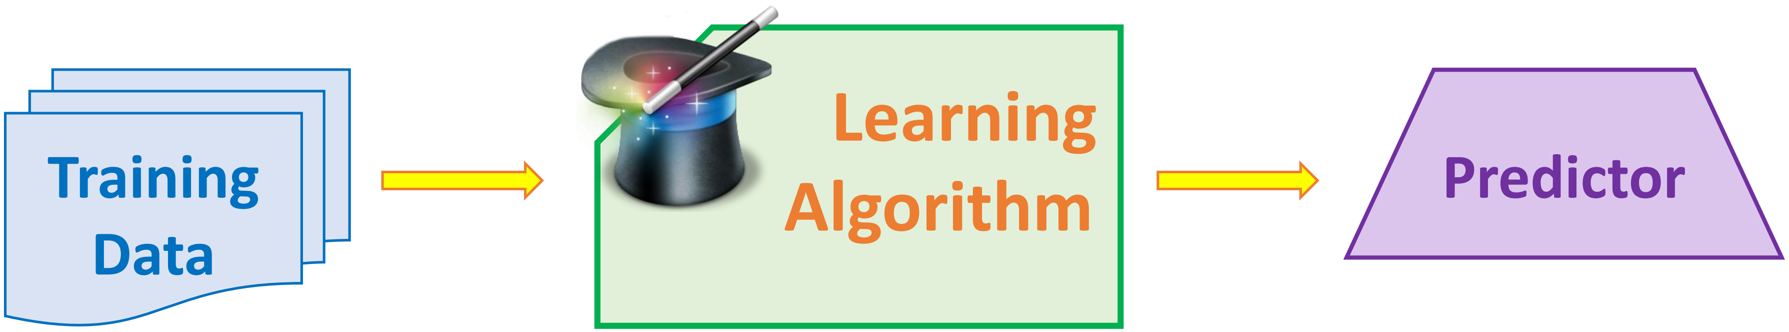

pinterest.com

Traditional “Batch” Learning

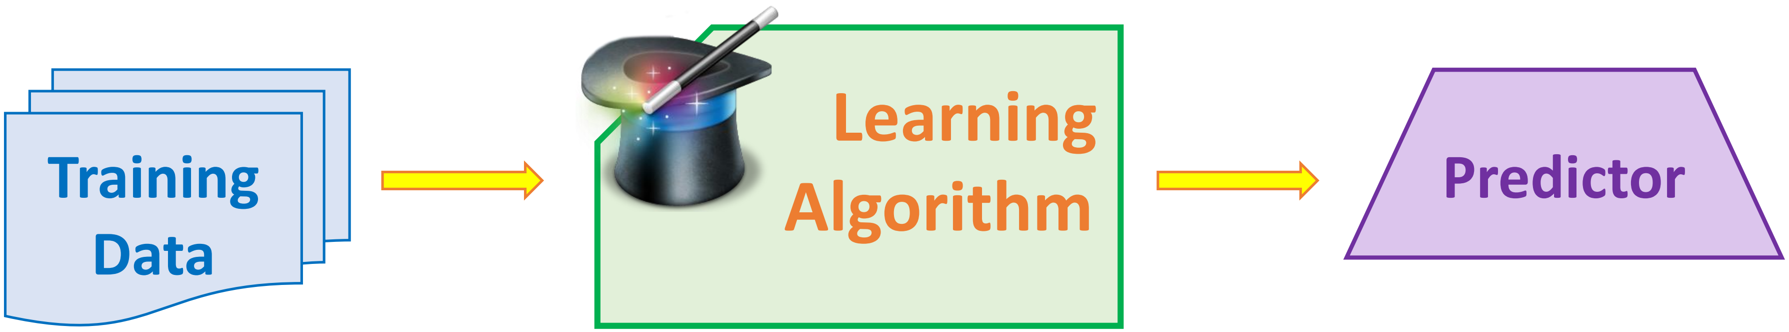

- A one round game between teacher and learner

Traditional “Batch” Learning

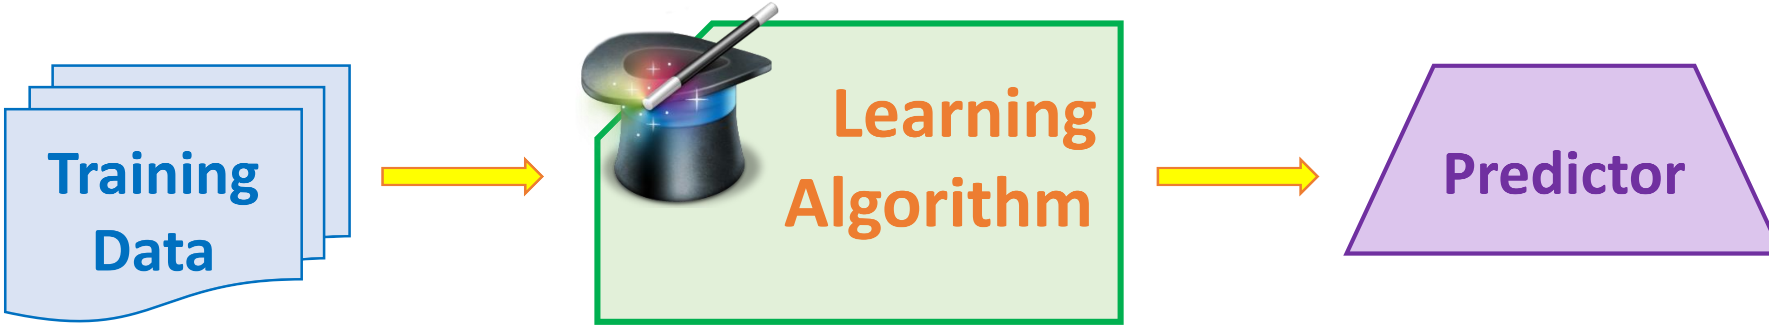

- A one round game between teacher and learner

pinterest.com

Traditional “Batch” Learning

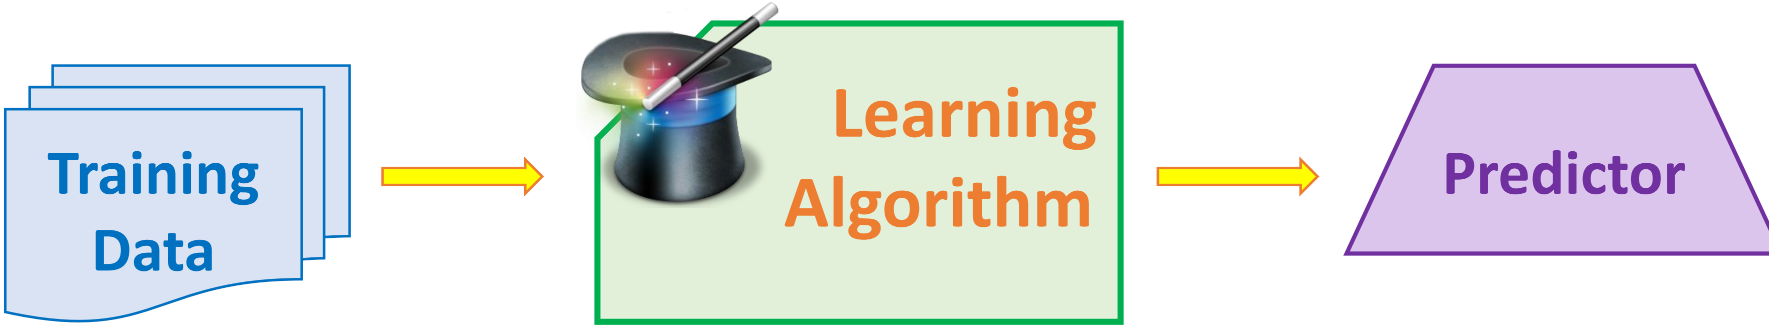

- A one round game between teacher and learner

pinterest.com

Traditional "Batch" Learning

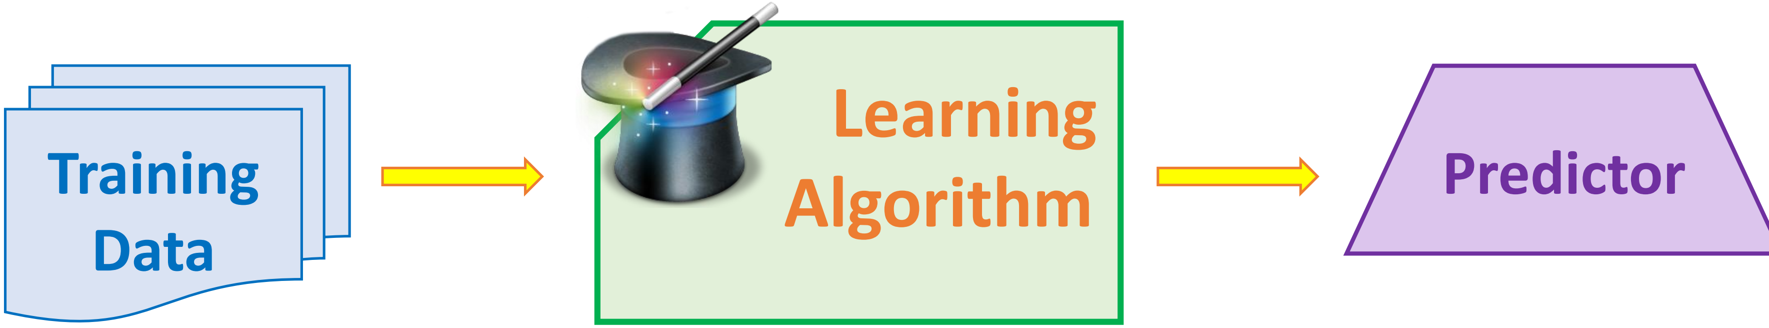

- A one round game between teacher and learner

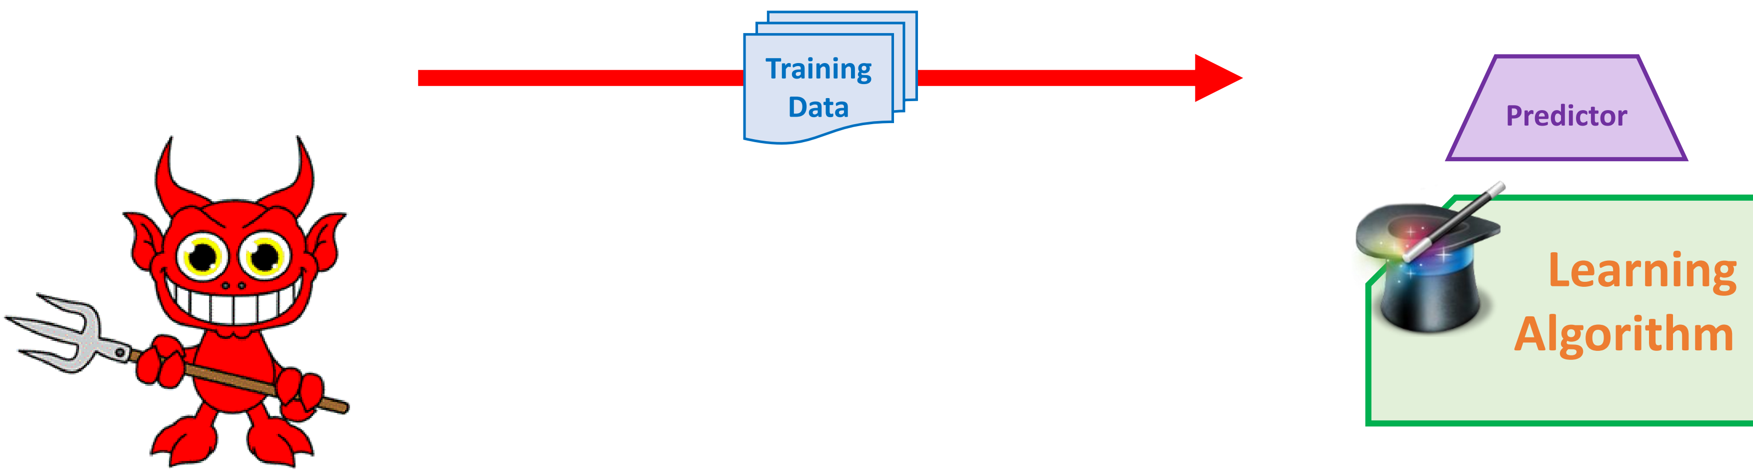

pinterest.com

Traditional “Batch” Learning

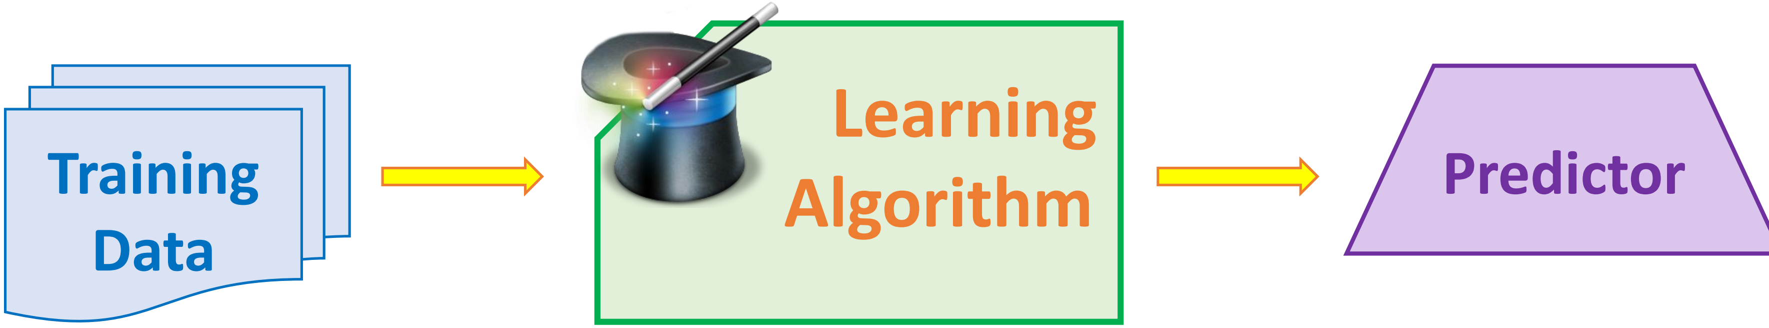

- A one round game between teacher and learner

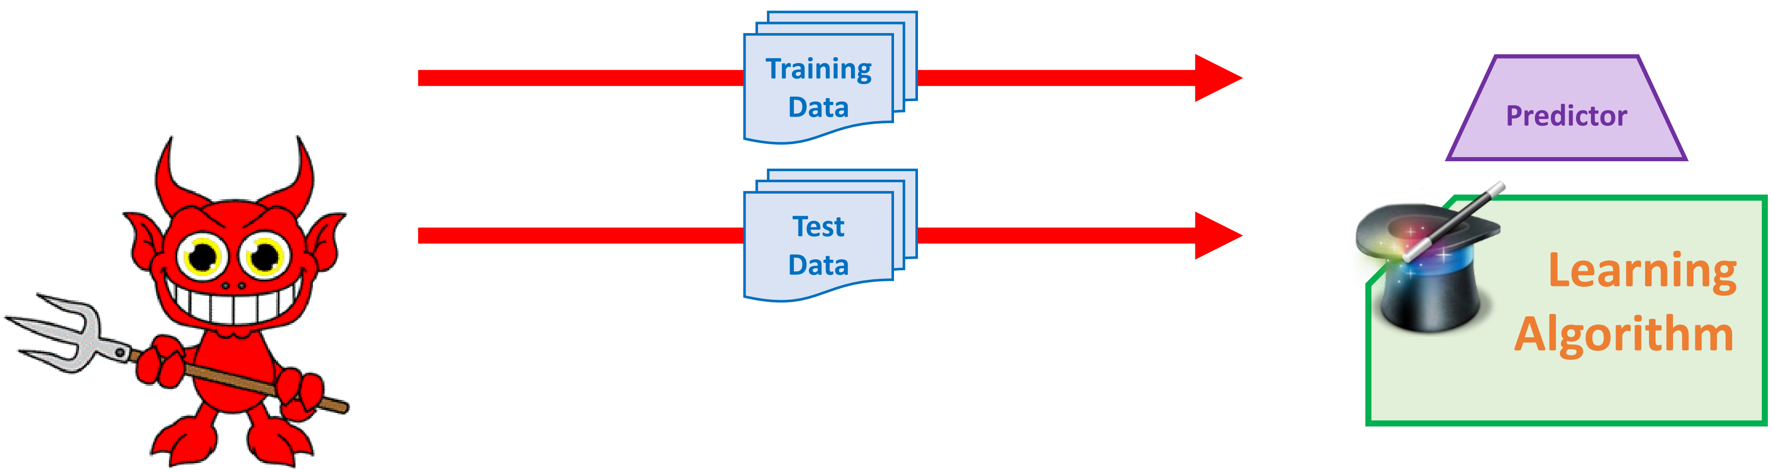

pinterest.com

Traditional "Batch" Learning

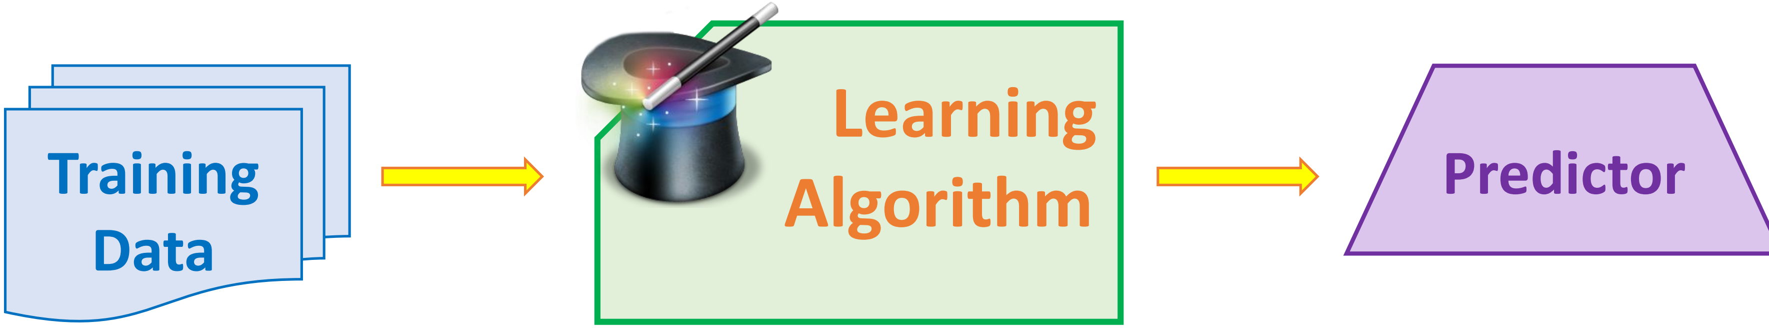

- A one round game between teacher and learner

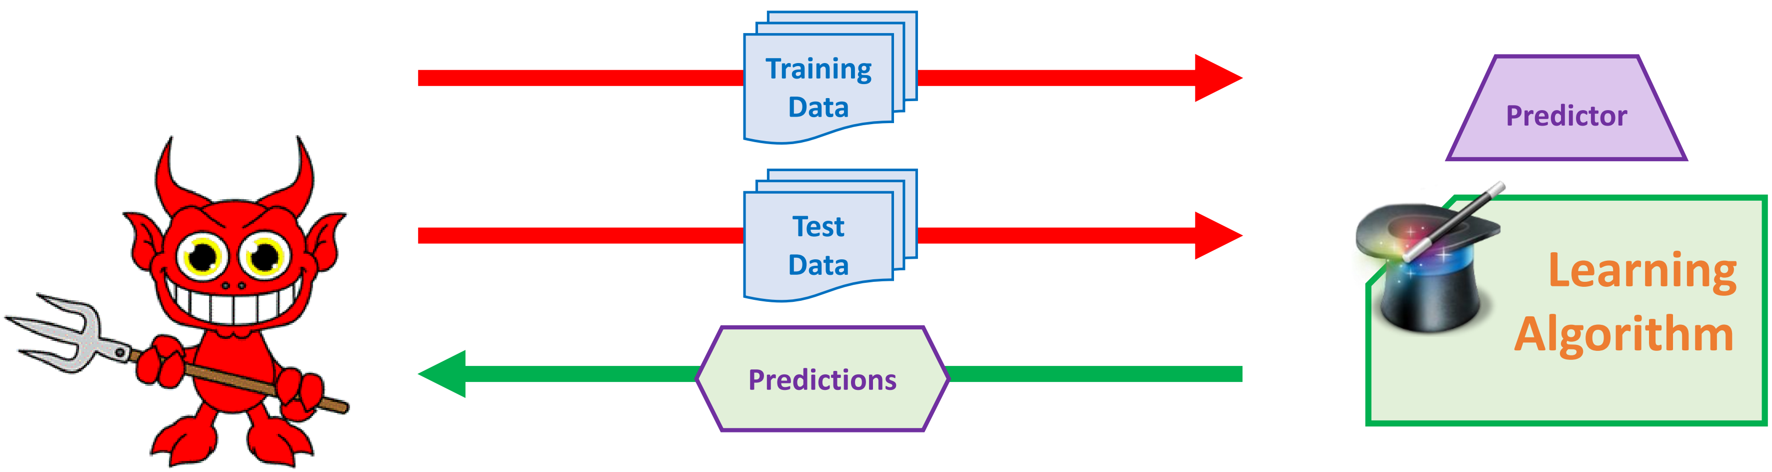

Traditional “Batch” Learning

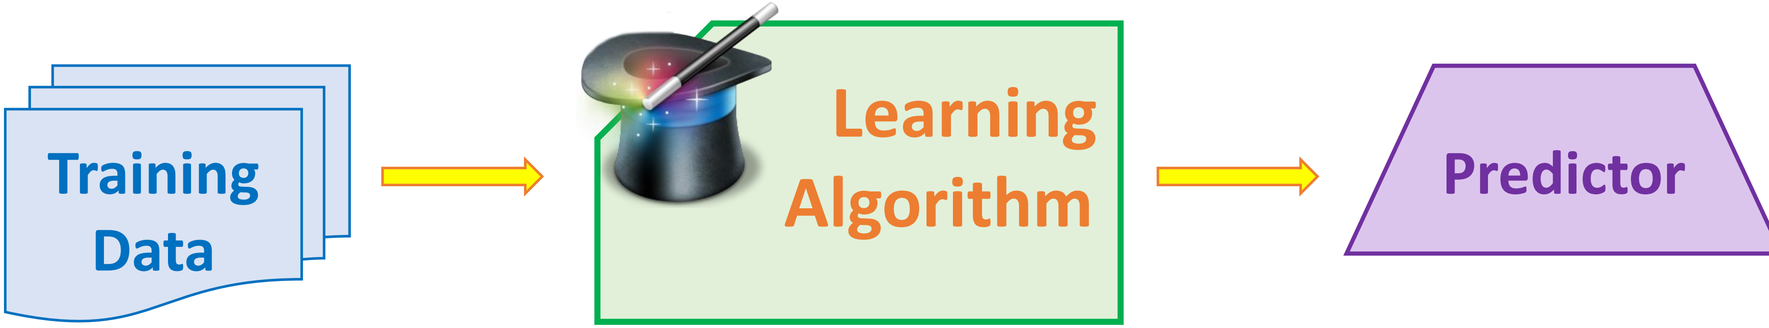

- A one round game between teacher and learner

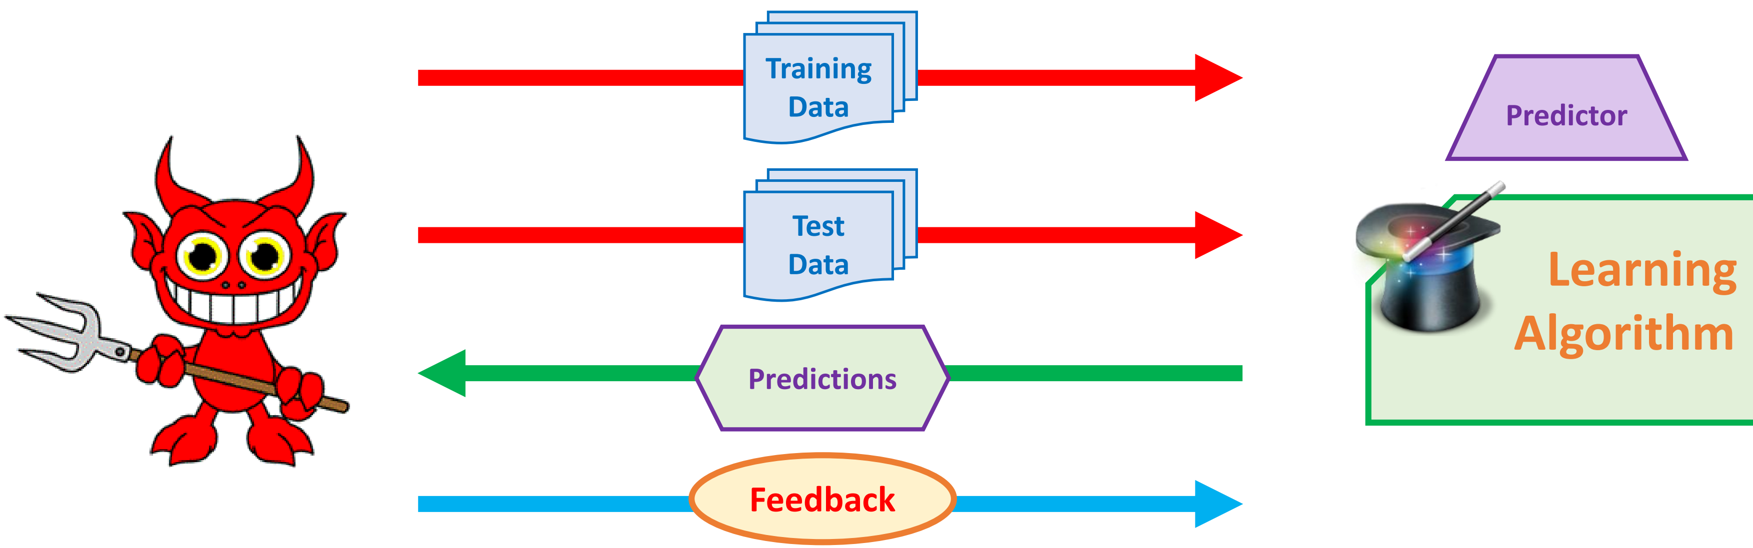

Traditional “Batch” Learning

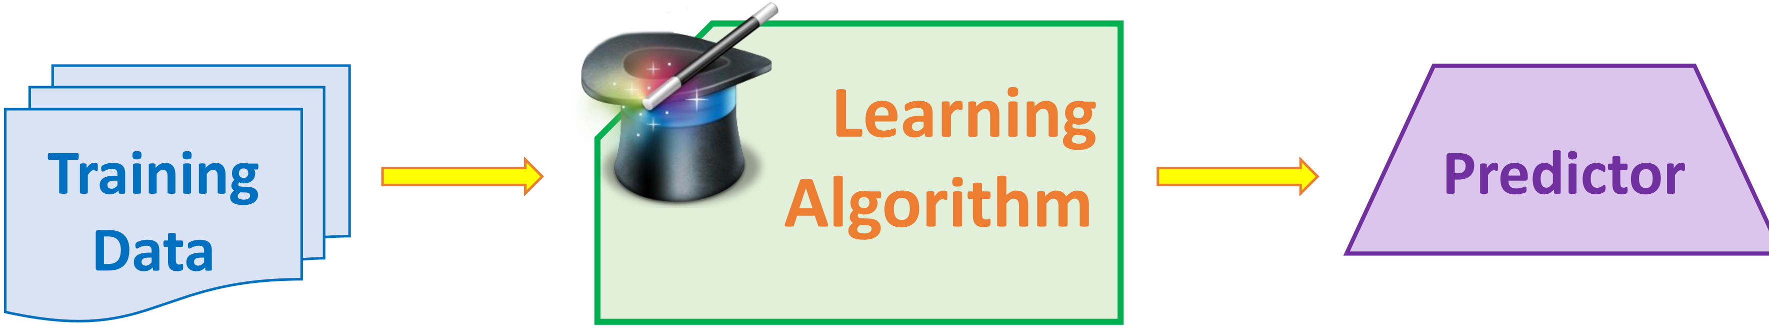

- A one round game between teacher and learner

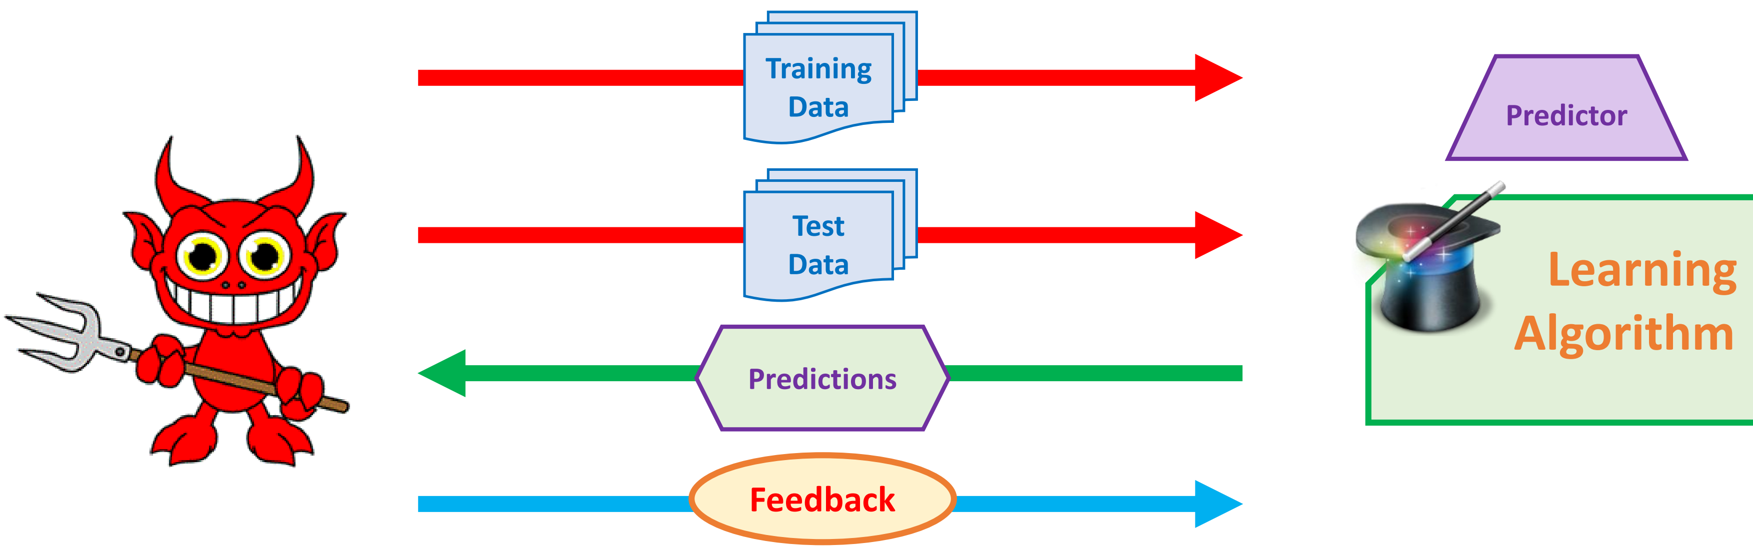

- As expected, each tries to outdo the other

pinterest.com

Online Learning

Online Learning

"The art and science of designing algorithms that can adapt to sequential data"

Online Learning

"The art and science of designing algorithms that can adapt to sequential data"

Learning Algorithm

Online Learning

"The art and science of designing algorithms that can adapt to sequential data"

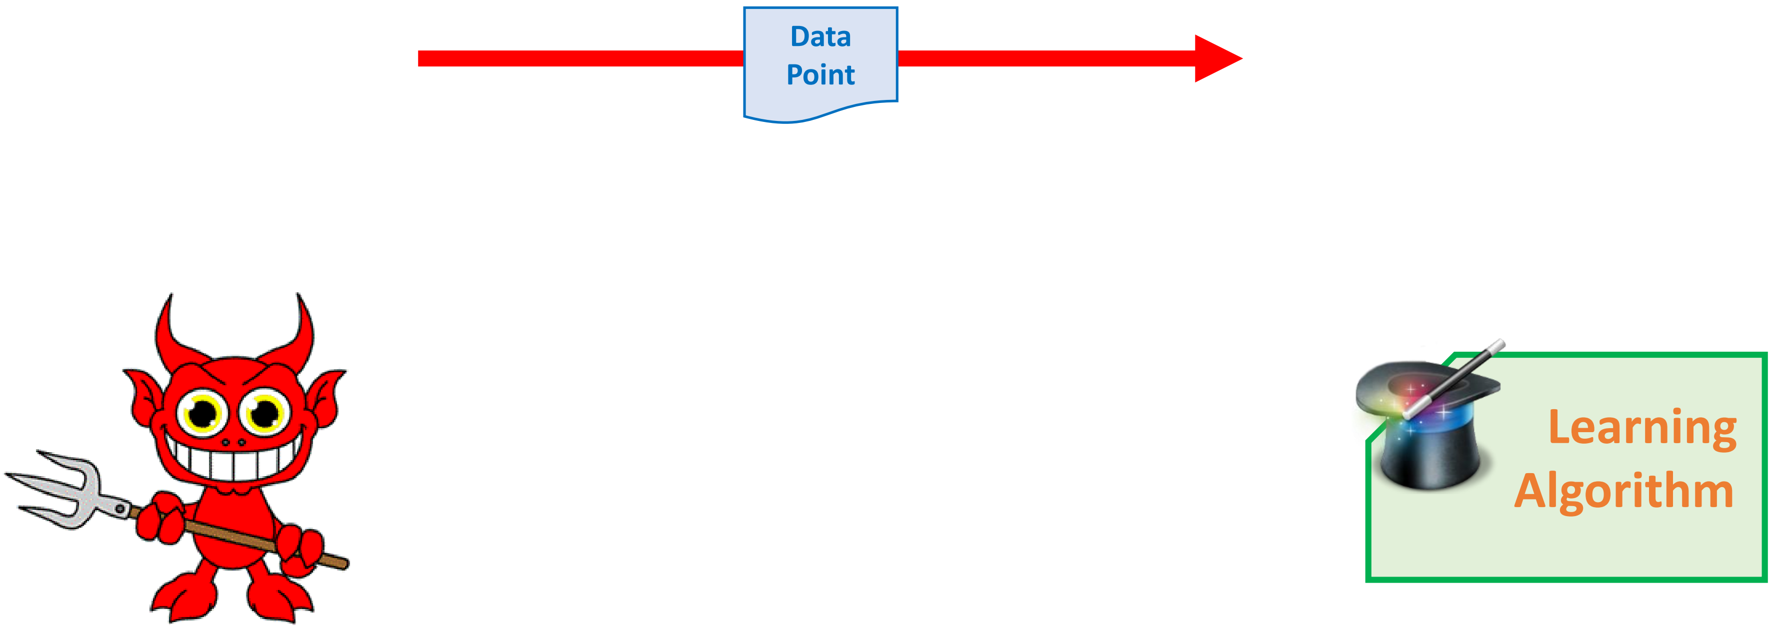

Online Learning

“The art and science of designing algorithms that can adapt to sequential data”

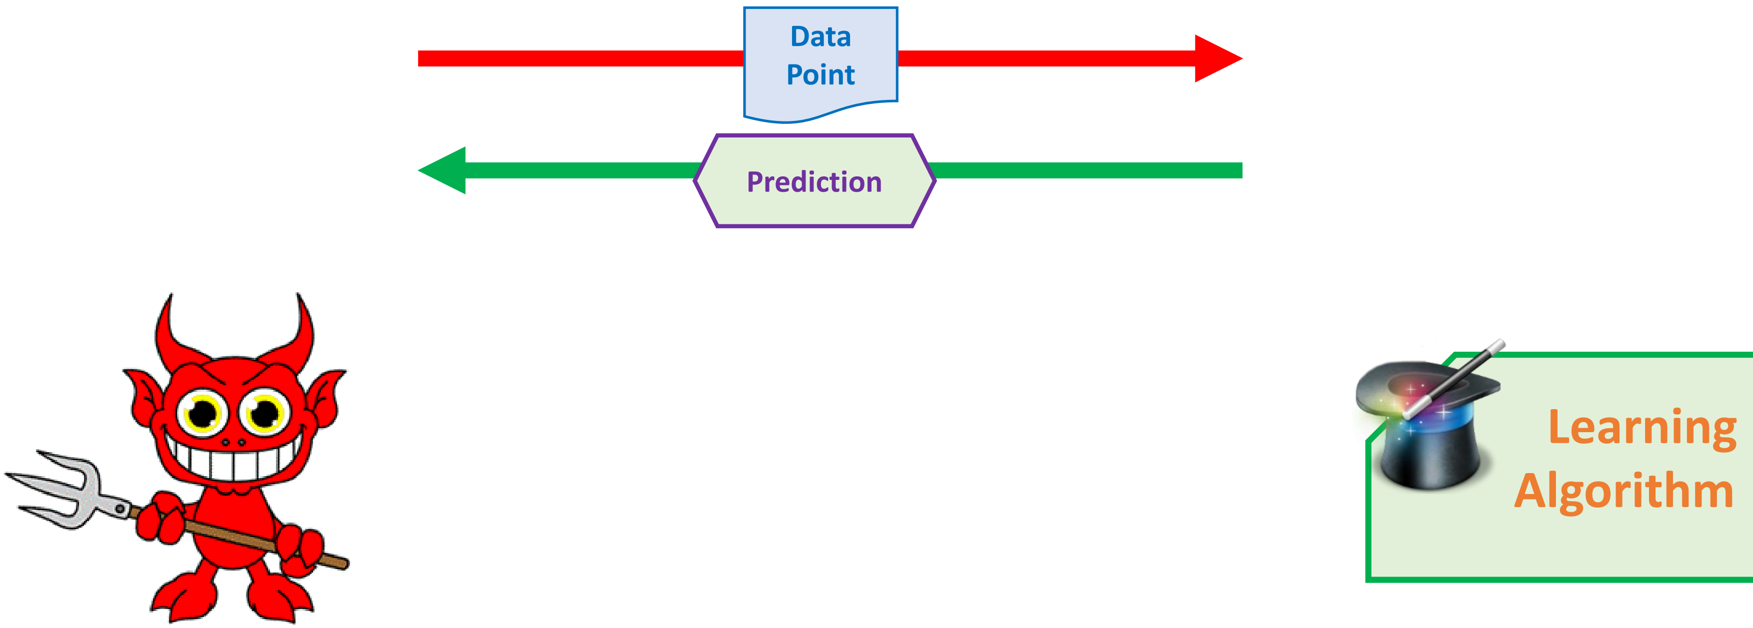

Online Learning

"The art and science of designing algorithms that can adapt to sequential data"

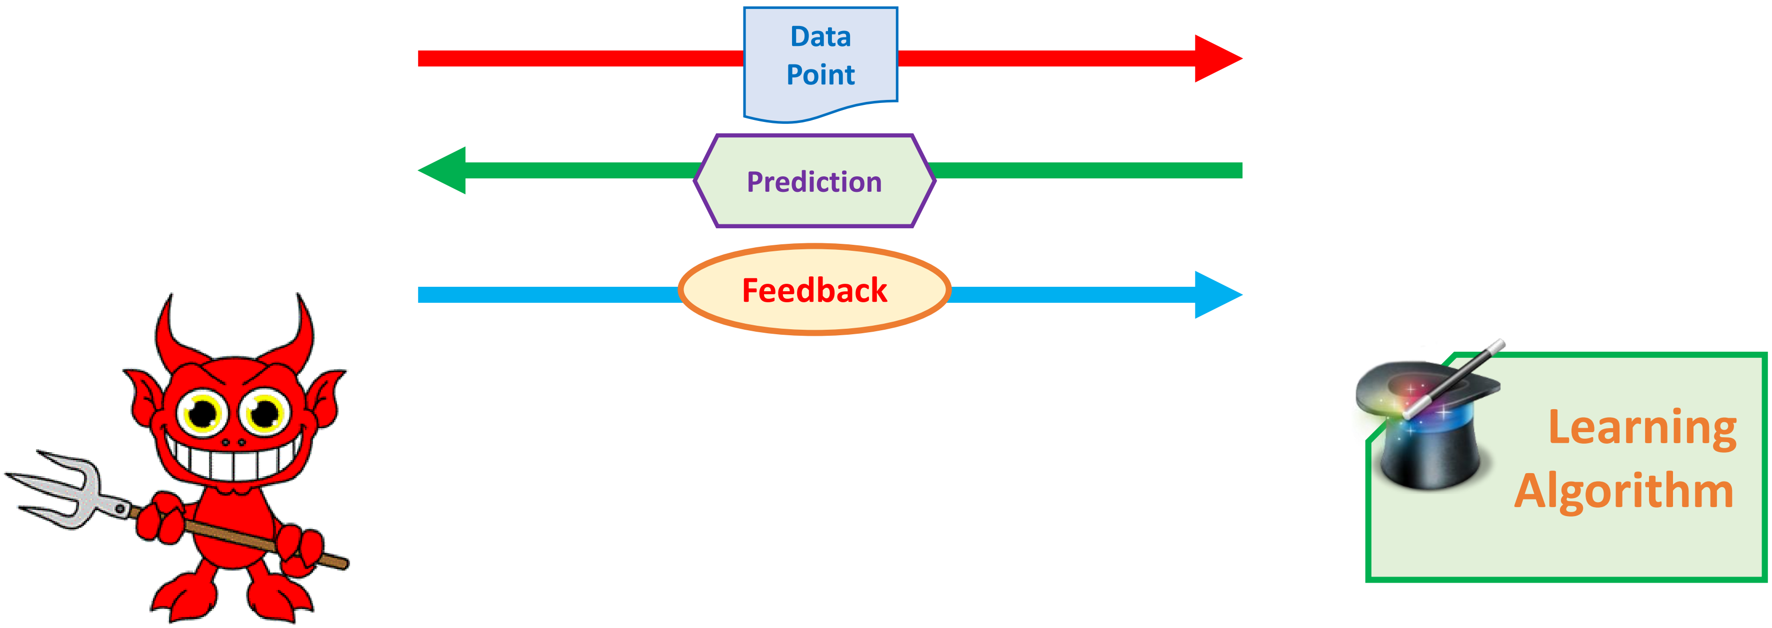

Online Learning

“The art and science of designing algorithms that can adapt to sequential data”

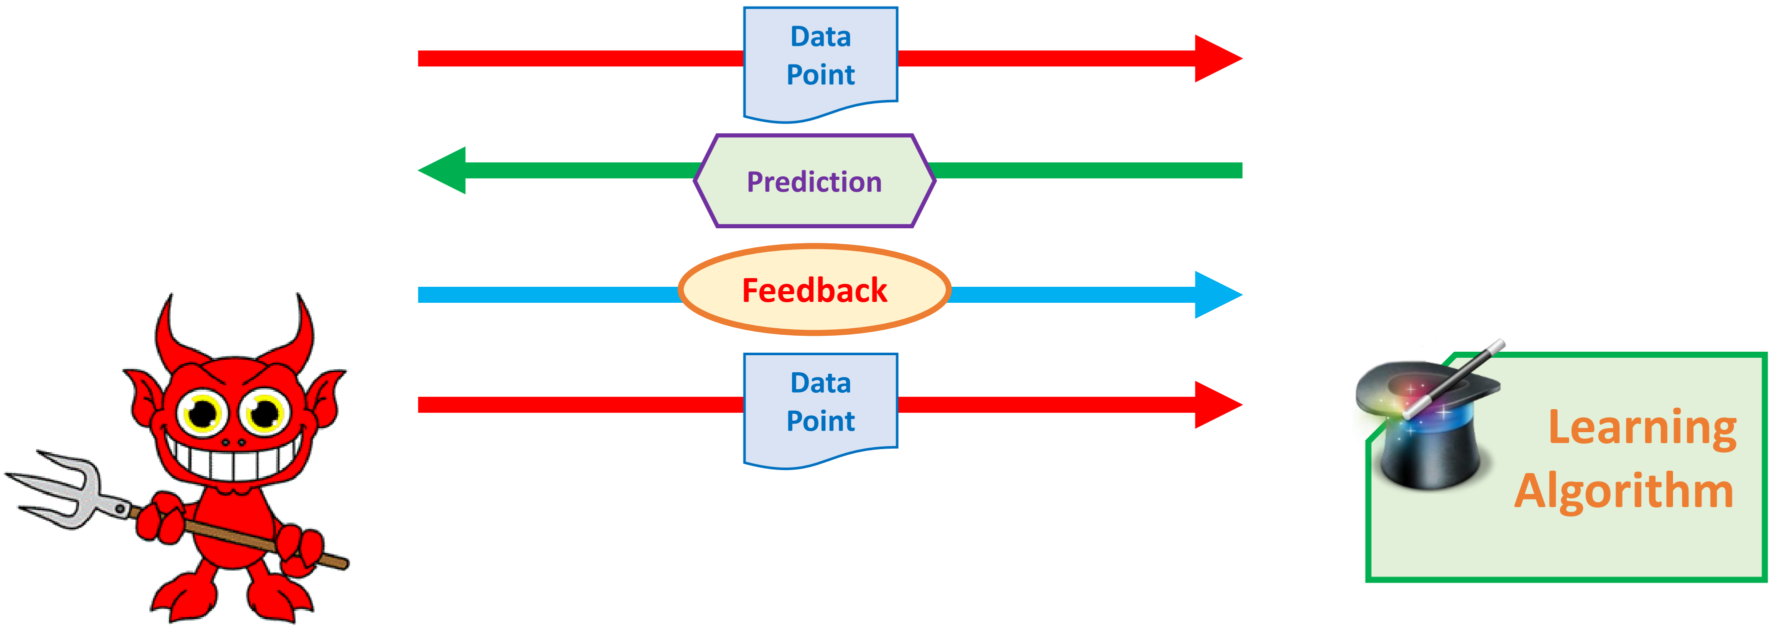

Online Learning

"The art and science of designing algorithms that can adapt to sequential data"

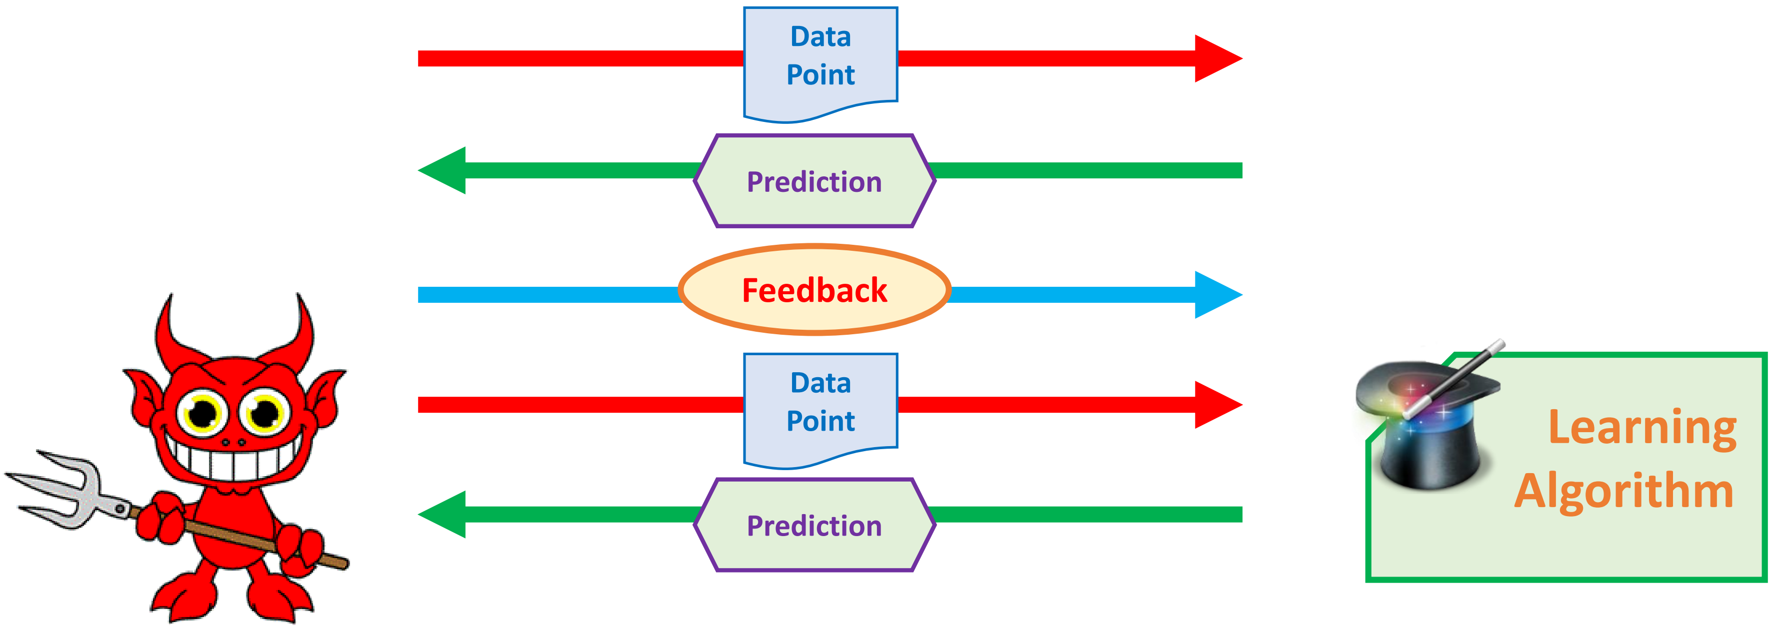

Online Learning

"The art and science of designing algorithms that can adapt to sequential data"

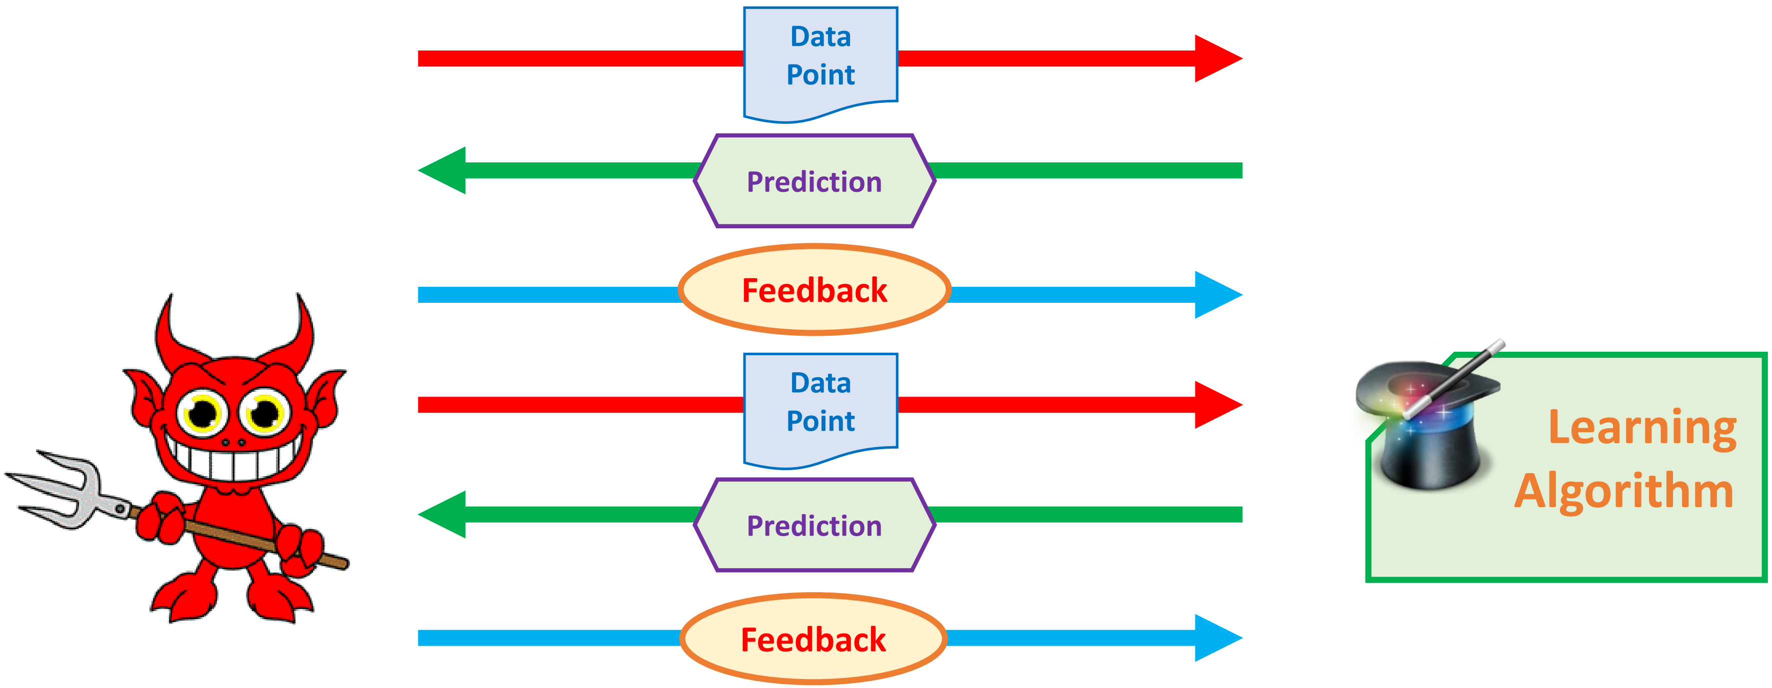

Online Learning

"The art and science of designing algorithms that can adapt to sequential data"

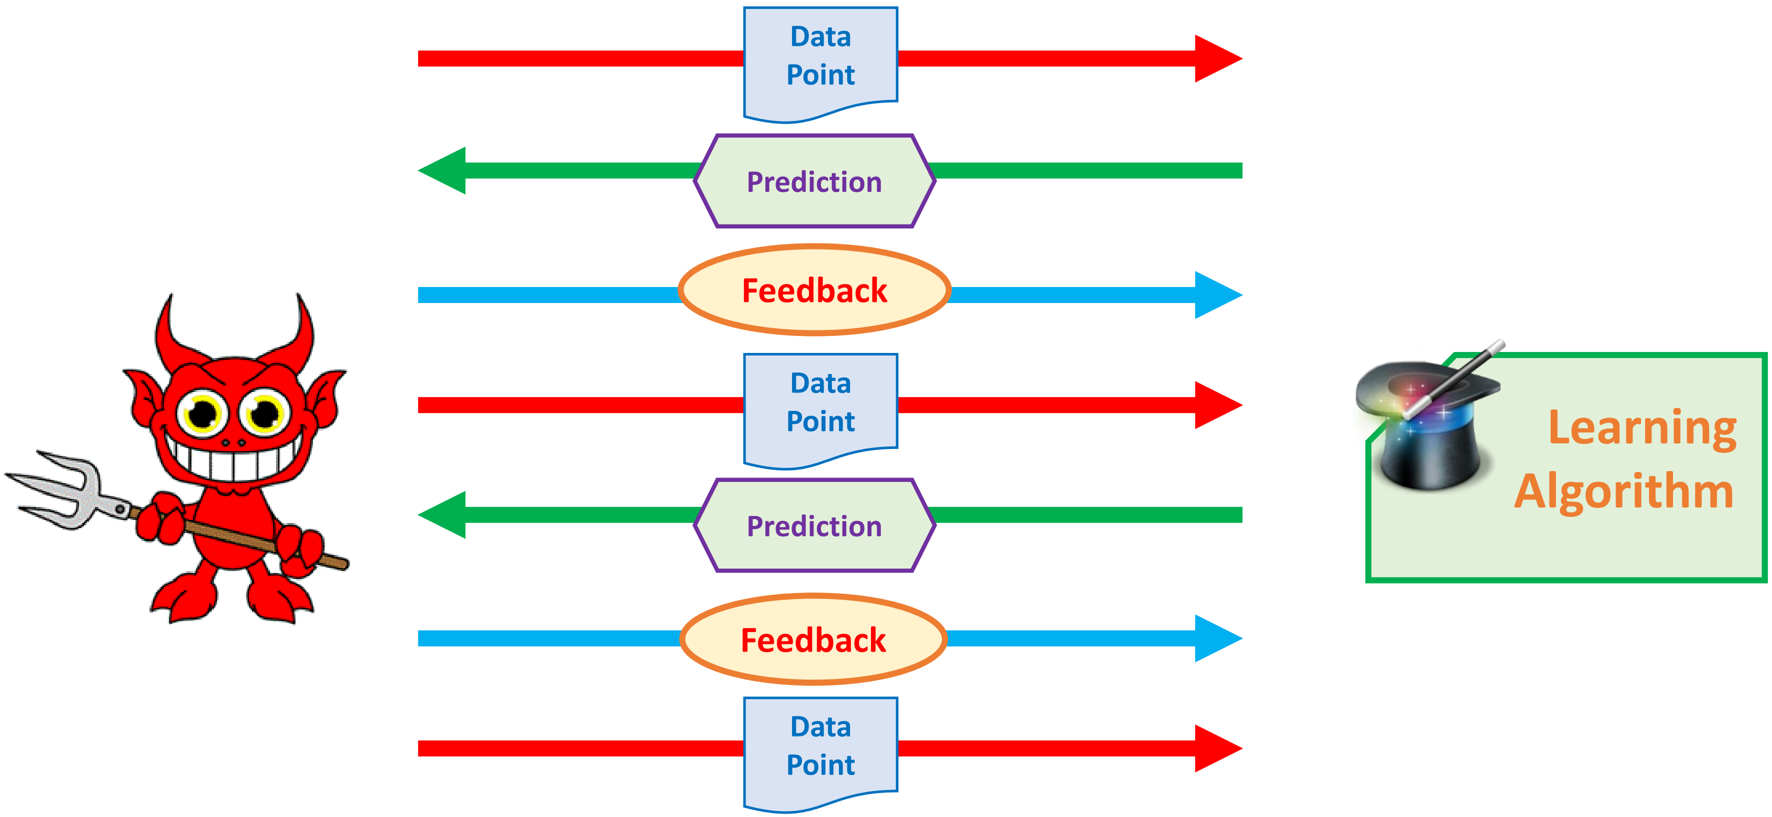

Online Learning

"The art and science of designing algorithms that can adapt to sequential data"

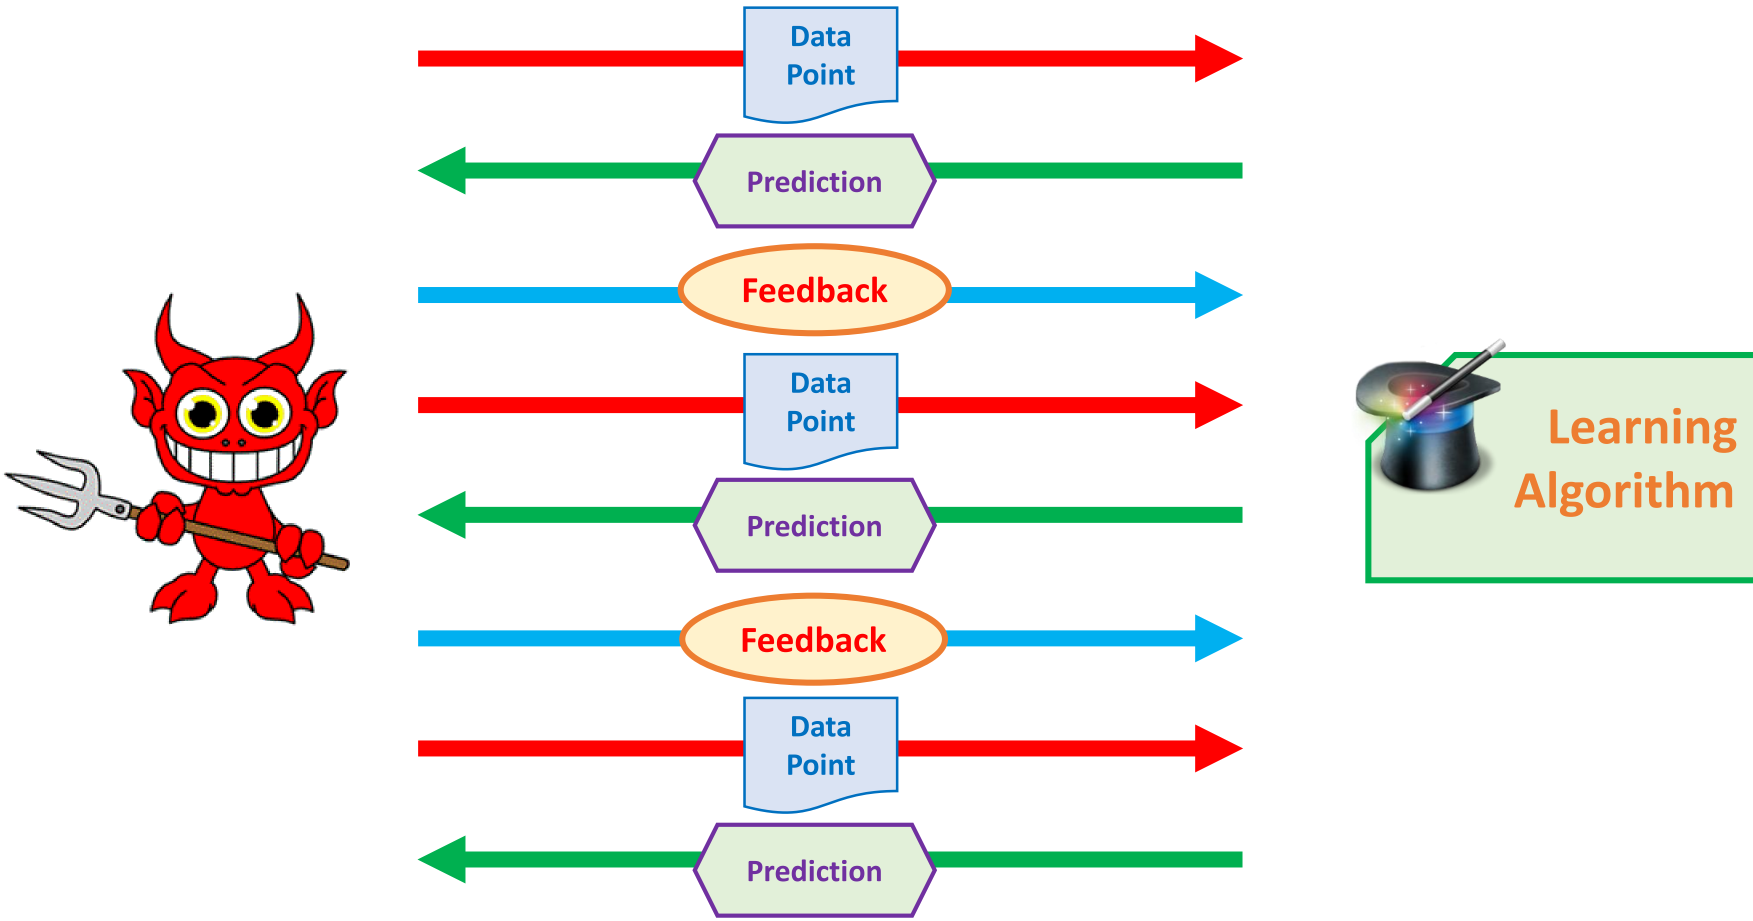

Online Learning

"The art and science of designing algorithms that can adapt to sequential data"

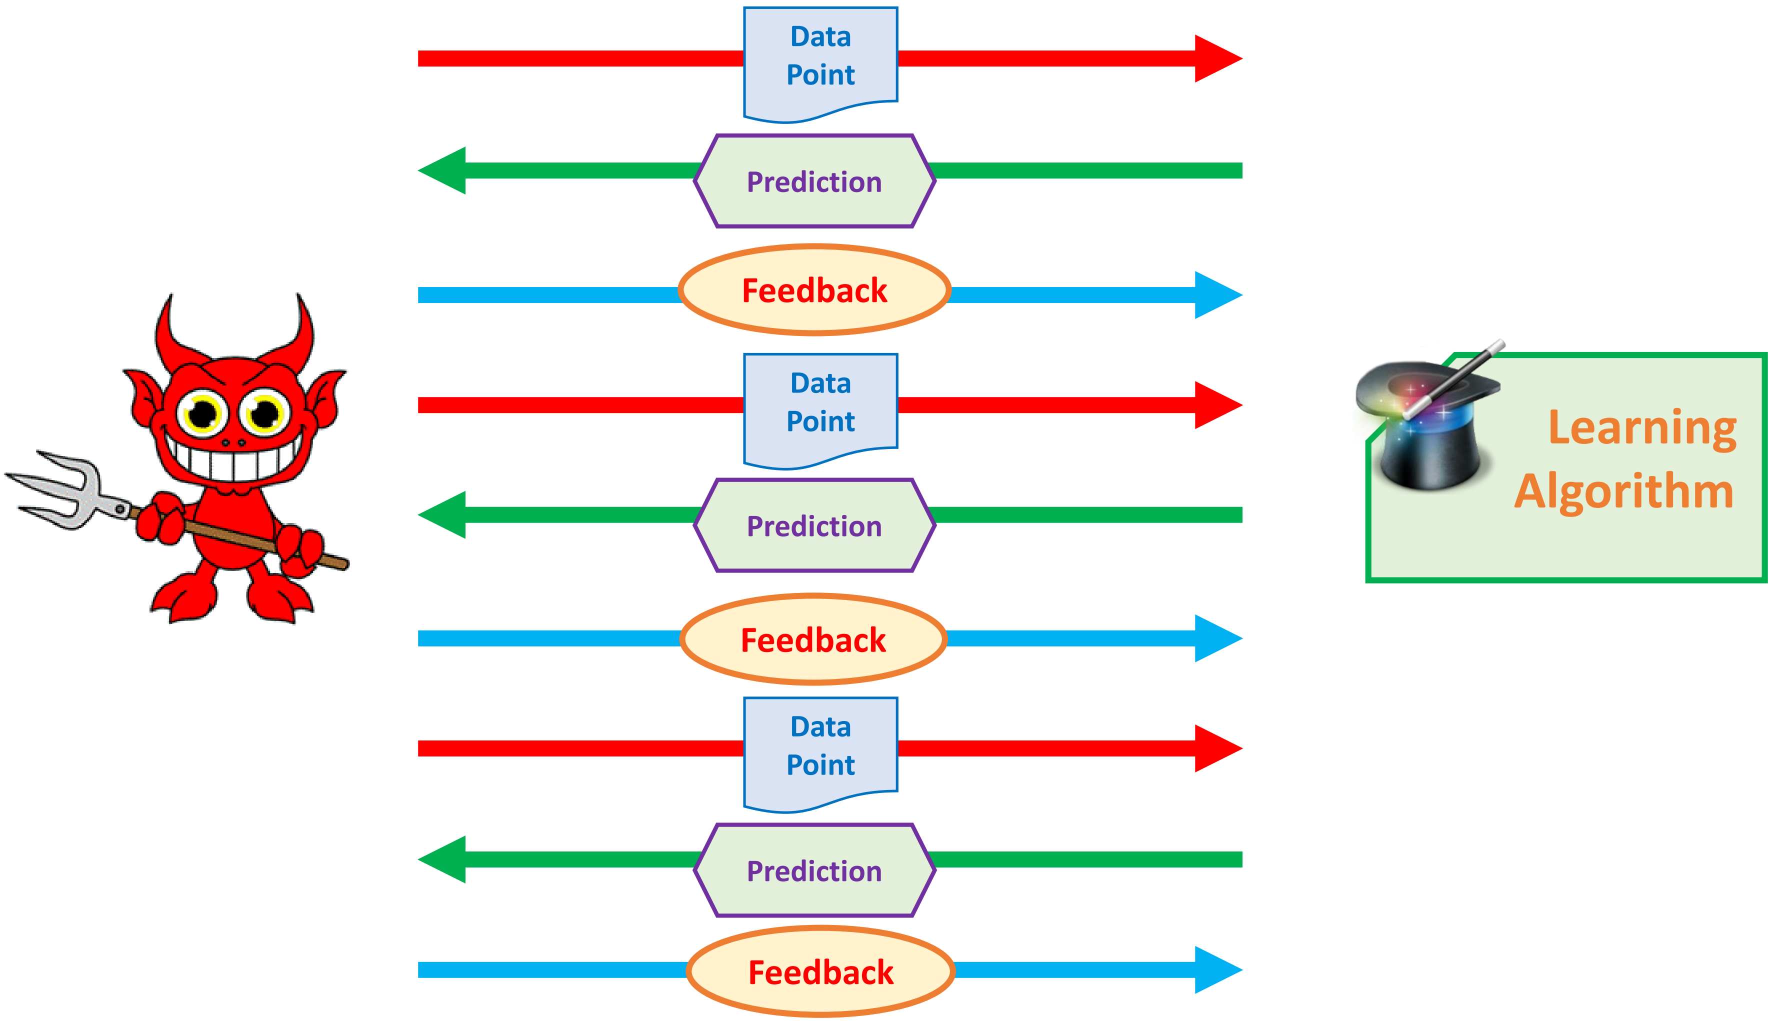

Online Learning

"The art and science of designing algorithms that can adapt to sequential data"

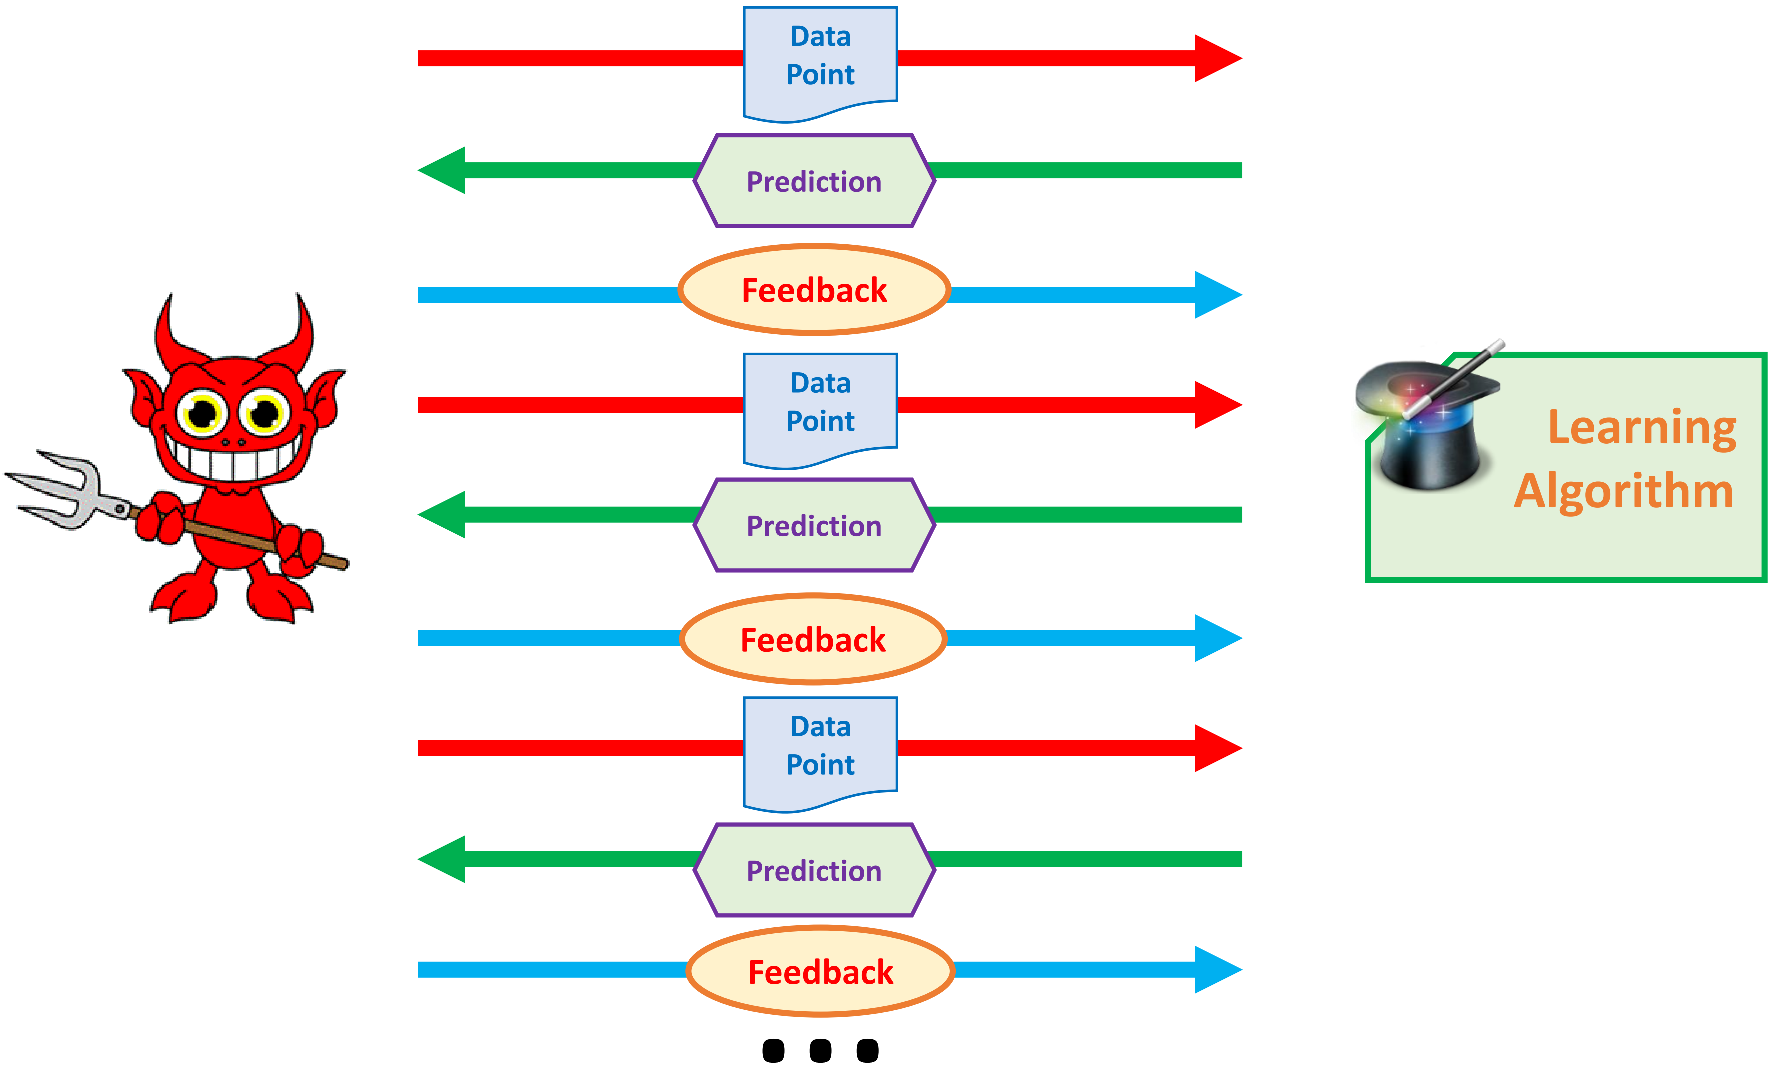

Online Learning

"The art and science of designing algorithms that can adapt to sequential data"

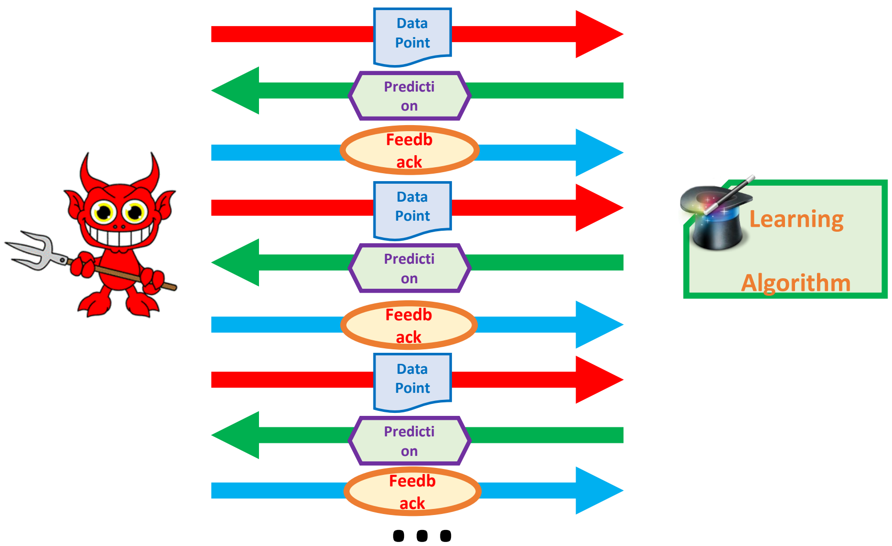

Online Learning

"The art and science of designing algorithms that can adapt to sequential data"

- Binary predictions – online classification

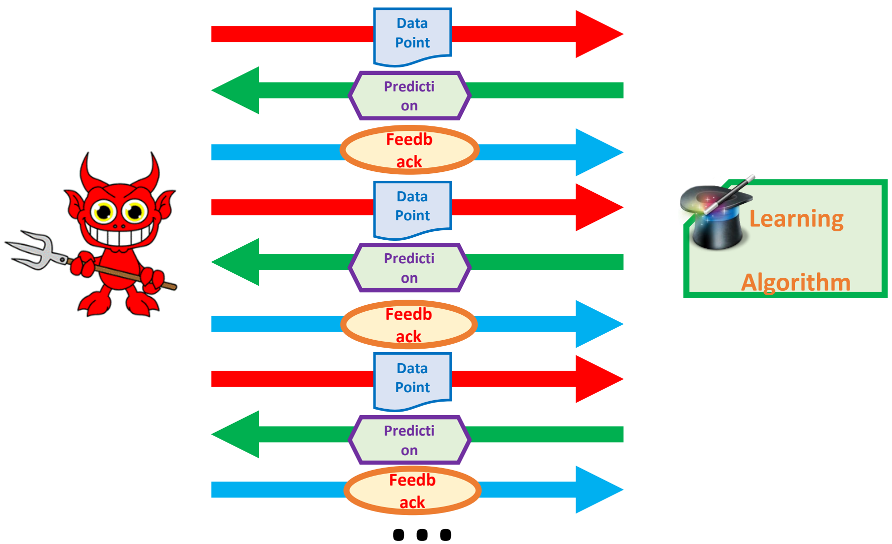

Online Learning

“The art and science of designing algorithms that can adapt to sequential data”

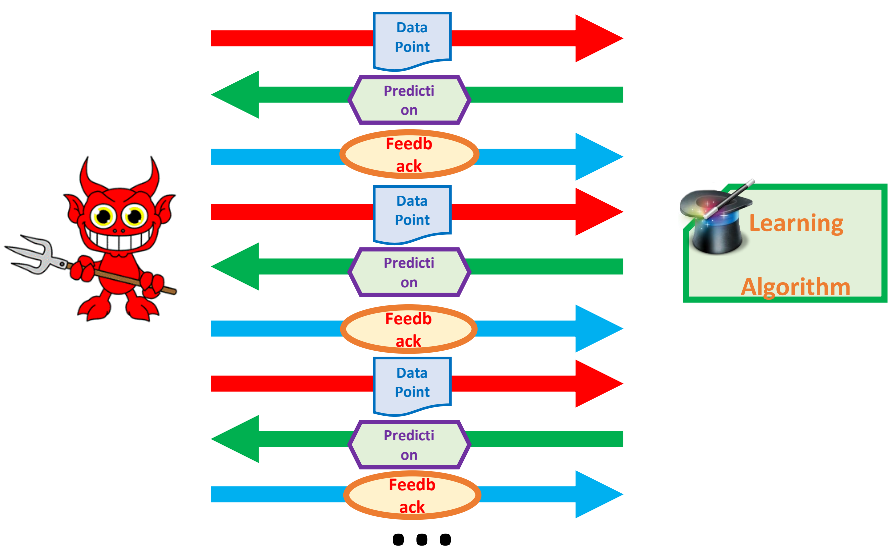

- Binary predictions – online classification
- Real predictions – online regression

Online Learning

"The art and science of designing algorithms that can adapt to sequential data"

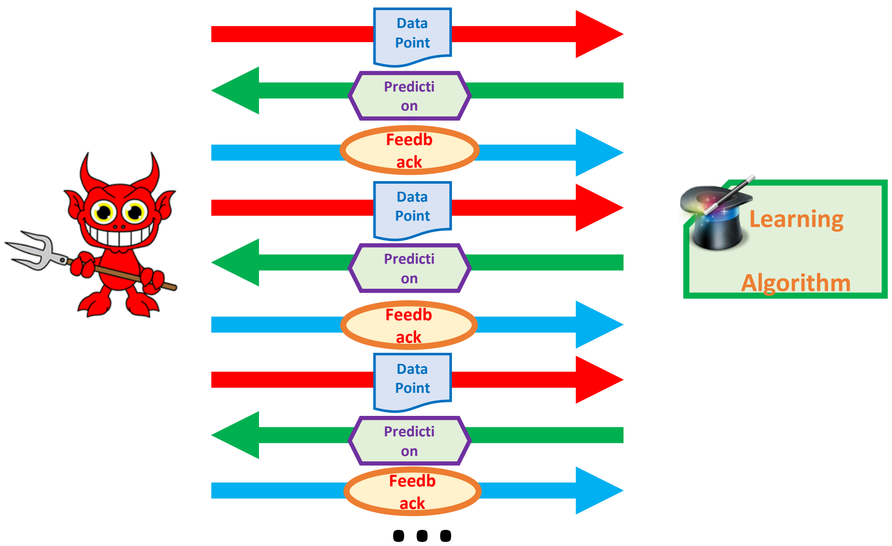

- Binary predictions – online classification
- Real predictions – online regression
- "Incomplete data" – bandit learning

Online Learning

"The art and science of designing algorithms that can adapt to sequential data"

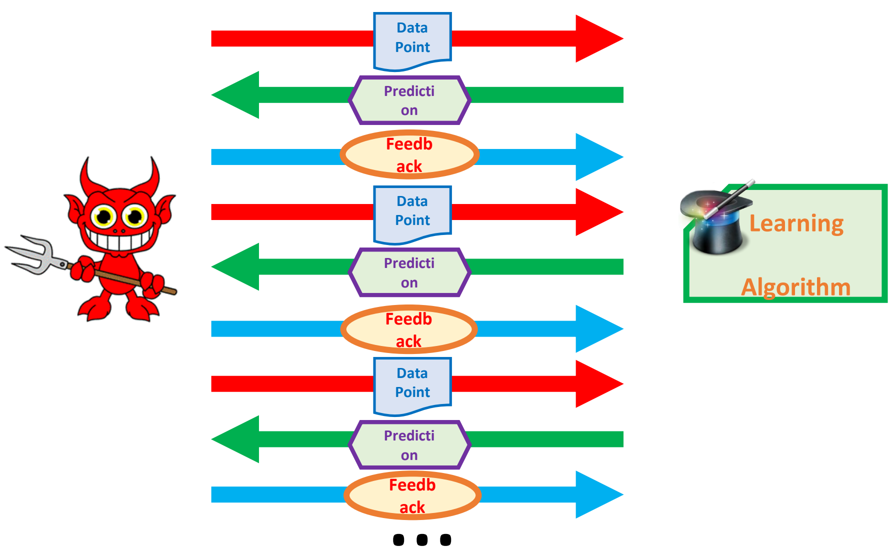

- Binary predictions – online classification
- Real predictions – online regression
- "Incomplete data" – bandit learning
- General model
  Data point: state
  Prediction: action
  Reinforcement Learning!!

Online Learning – Applications

- Online spam filtering
  - Data point: email description
  - Prediction: spam/okay
  - Feedback: Correctness of prediction

- Portfolio selection
  - Data point: market description
  - Prediction: investment profile
  - Feedback: revenue earned/lost

- Recommendation systems
  - Data: user profile
  - Prediction: items to buy/movies to watch
  - Feedback: click on suggested items

NETFLIX flipkart.com amazon

mediasuite.ca, touchpointsa.com, cordcuttersnews.com, phx.corporate-ir.net, rednewswire.com

Online Learning – Applications

- Ad-placement systems
  - Data: user profile, history
  - Prediction: ads displayed
  - Feedback: click, purchase

- Weather prediction
  - Data: Recent met data, historical
  - Prediction: rain, amount
  - Feedback: actual weather

- Stock price prediction
  - Data: market description, past prices
  - Prediction: future prices
  - Feedback: actual prices

techgyd.com, metoffice.gov.uk, cnbc.com

Online Optimization

$$\min_{\mathbf{x} \in \Omega} f(\mathbf{x}) = \sum_{i=1}^{n} g(\mathbf{x}; \theta_i) + r(\mathbf{x})$$

- Cousin of online learning - optimizing over data streams
- Immensely useful in optimization over large datasets
- Extends traditional "batch" optimization methods
  - Gradient descent, Mirrored descent
  - Newton's method
- Widely used method
  - defacto standard in several areas
  - SVM solvers – LibSVM, Liblinear
  - Training deep nets

Course Contents

- **Online Prediction with Full Feedback**
  - Online classification, regression
  - Learning with expert advice, portfolio selection
- **Online Convex Optimization**
  - Review of batch optimization
  - FTRL, OGD, OMD, SGD (OMG right??)
- **Online Prediction with Limited Feedback**
  - Stochastic/adversarial multi-armed bandits
  - Linear and contextual bandits
- **Advanced topics***
  - SVRG, Minimax rates, Zero-order optimization, shifting experts
- Feedback on topics appreciated

How to Feel no Regret

Online Classification

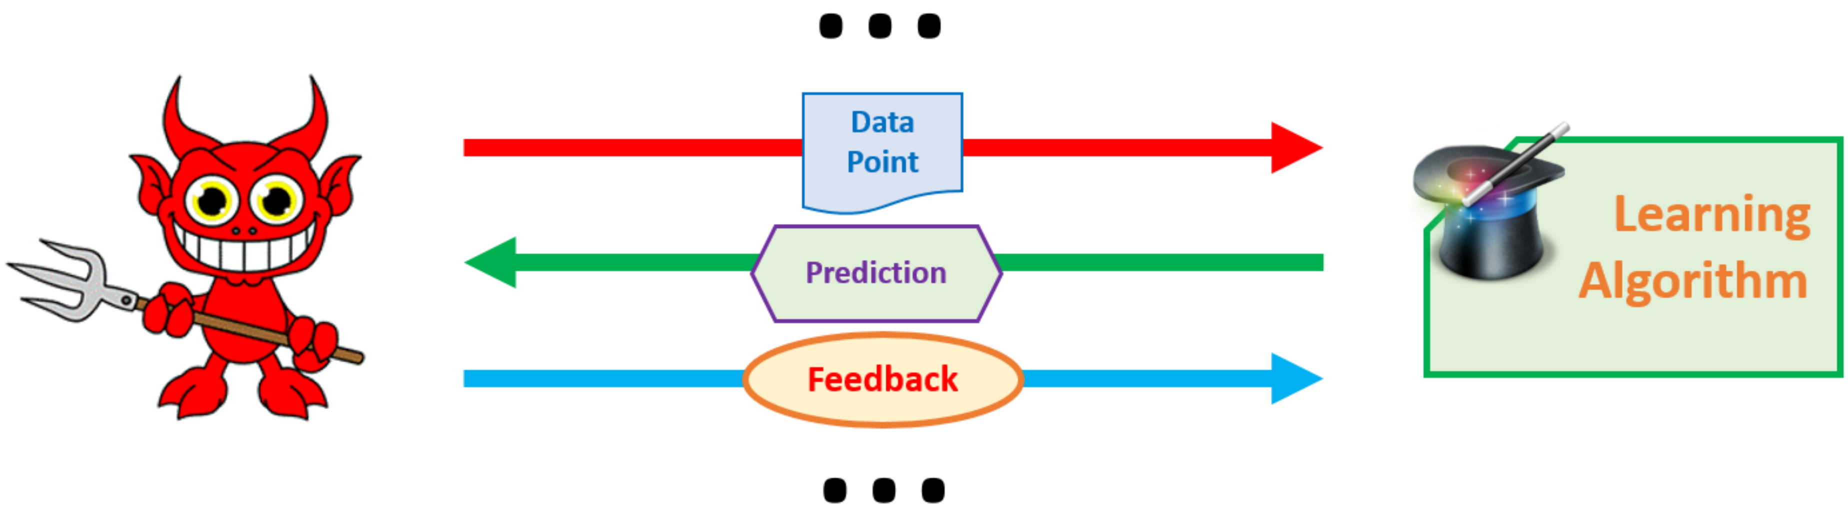

- At each time step t
  - Learner receives a *context* $\mathbf{x}_t$
  - Learner proposes a label $\hat{y}_t \in \mathcal{Y}$
  - Teacher provides true label as feedback $y_t \in \mathcal{Y}$
  - Learner incurs a loss of $\ell(\hat{y}_t, y_t)$
    - Example: $\ell^{0/1}(\hat{y}_t, y_t) = \mathbb{I}\{\hat{y}_t \neq y_t\}$
- Mistake Bound
  - A bound on the quantity $\sum_{t=1}^{T} \ell(\hat{y}_t, y_t)$

Online Learning

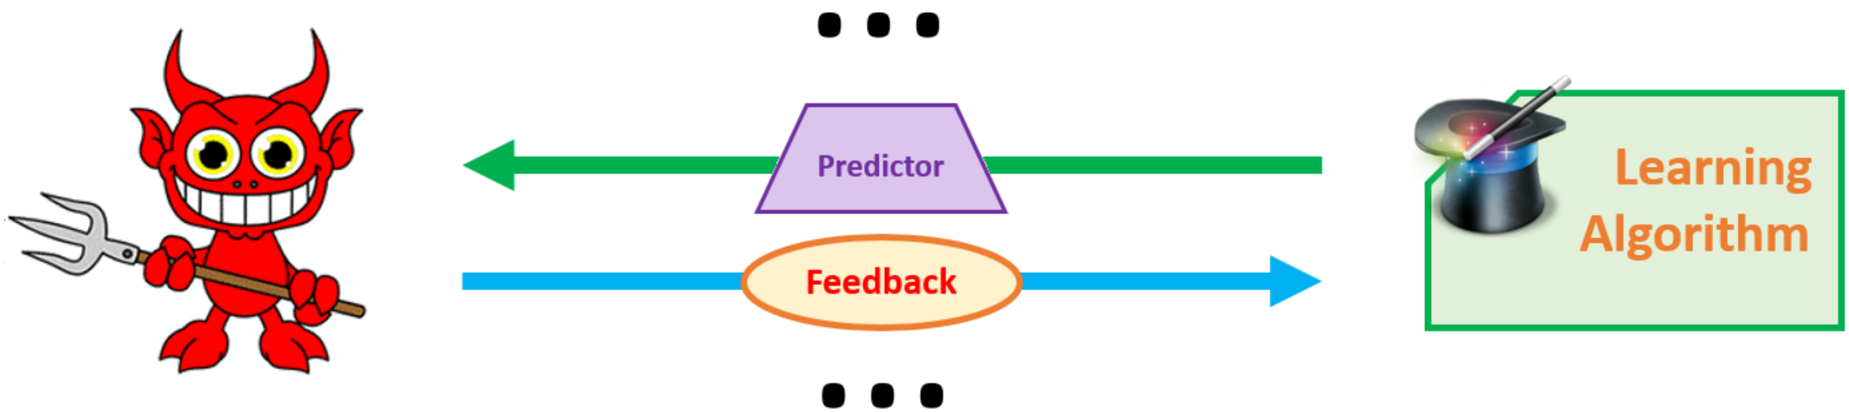

- At each time step t
  - Learner proposes a predictor $\mathbf{w}_t$
  - Teacher provides a *penalty* function as feedback $\ell_t(\cdot)$
  - Learner incurs a penalty $\ell_t(\mathbf{w}_t)$
- Typically: $\ell_t(\mathbf{w}) = \ell(\mathbf{w}, (\mathbf{x}_t, y_t))$
- Online linear regression: $\ell_t(\mathbf{w}) = (\langle \mathbf{w}, \mathbf{x}_t \rangle - y_t)^2$
- Cumulative Penalty
  - A bound on the quantity $\sum_{t=1}^{T} \ell_t(\mathbf{w}_t)$

Some Points to Note

Some Points to Note

- The teacher chooses the true labels/penalty functions **after** the learner has made his move

Some Points to Note

- The teacher chooses the true labels/penalty functions **after** the learner has made his move
- Consistent with what happened in the "batch" mode

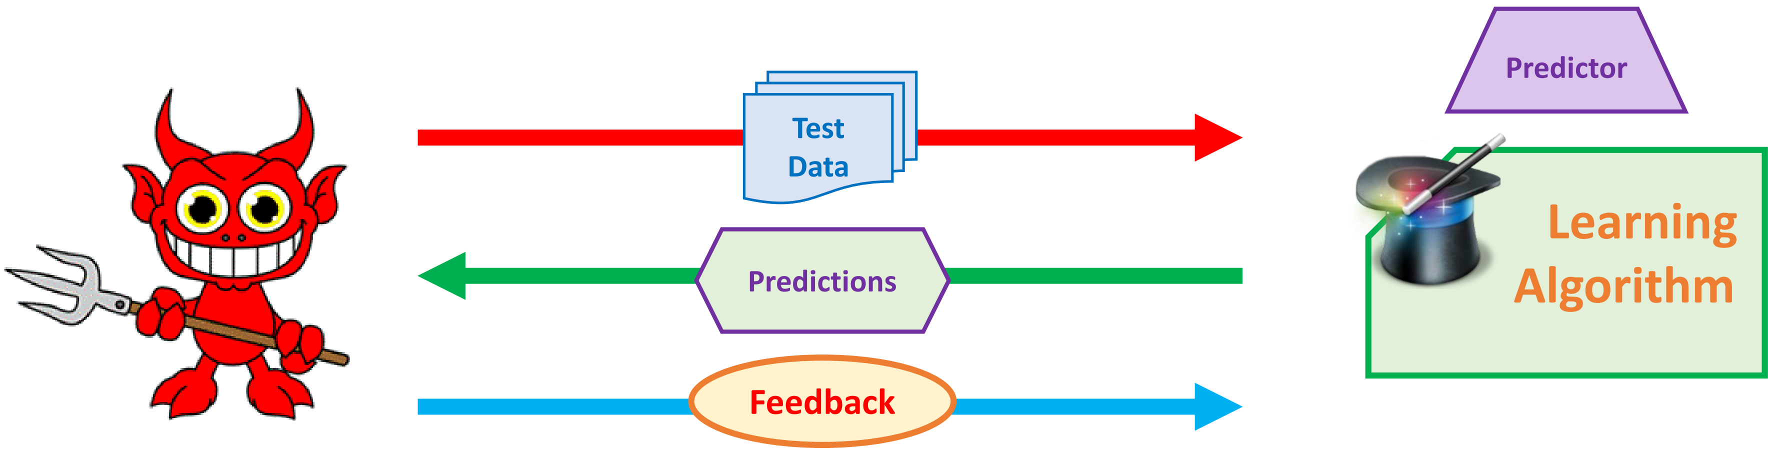

Some Points to Note

- The teacher chooses the true labels/penalty functions **after** the learner has made his move
- Consistent with what happened in the “batch” mode

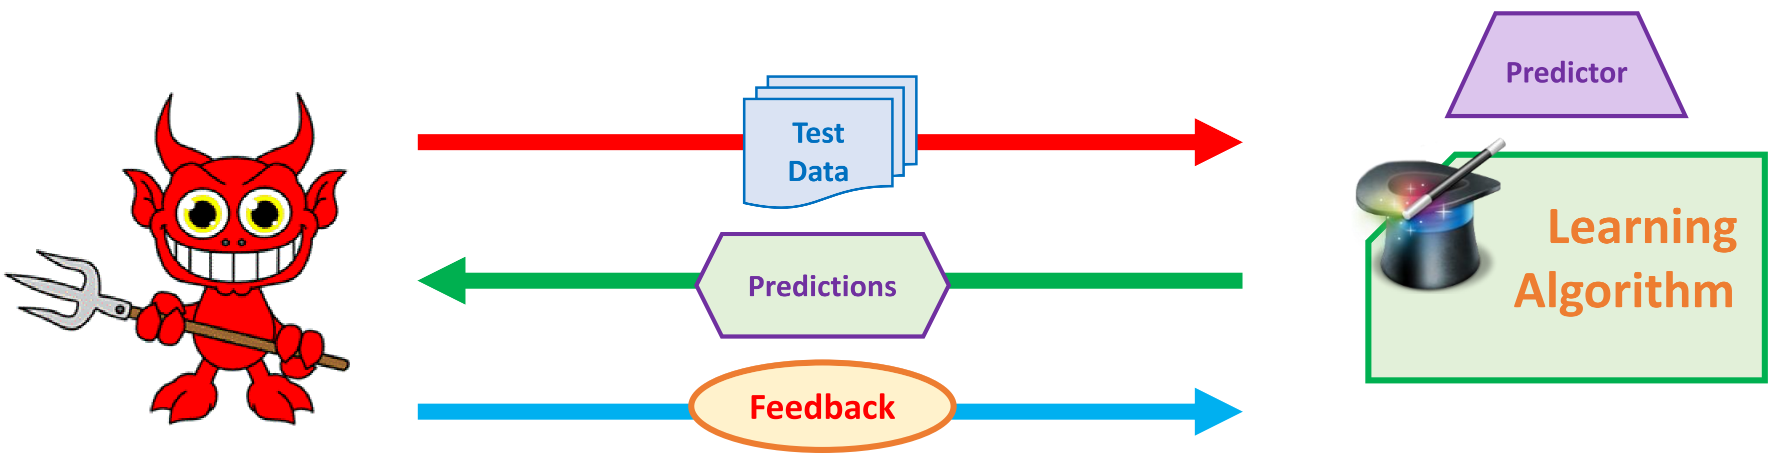

- How does the teacher generate feedback?
  - Stochastically $\ell_t(\cdot) = \ell(\cdot, (\mathbf{x}_t, y_t)), (\mathbf{x}_t, y_t) \sim \mathcal{D}$
  - Adversarially $\ell_t(\cdot) = \ell(\cdot, (\mathbf{x}_t, y_t))$

Some Points to Note

- The teacher chooses the true labels/penalty functions **after** the learner has made his move
- Consistent with what happened in the "batch" mode

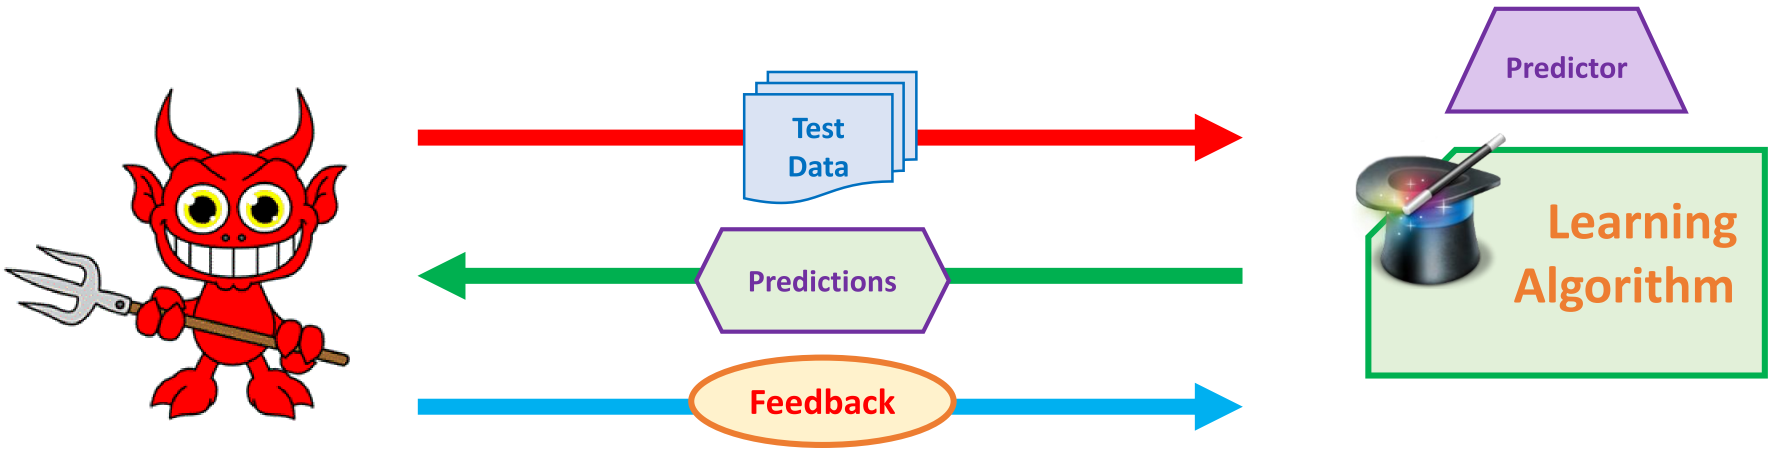

- How does the teacher generate feedback?
  - Stochastically $\ell_t(\cdot) = \ell(\cdot, (\mathbf{x}_t, y_t)), (\mathbf{x}_t, y_t) \sim \mathcal{D}$
  - Adversarially $\ell_t(\cdot) = \ell(\cdot, (\mathbf{x}_t, y_t))$
- How do we make sense of these settings?

Notion of Regret

Notion of Regret

- How do we distinguish between situations where
  - Data is easy and we should expect learner to do very well
  - Data is ridiculous and no learner can do well

Notion of Regret

- How do we distinguish between situations where
  - Data is easy and we should expect learner to do very well
  - Data is ridiculous and no learner can do well
- Mistake and Cumulative Penalties do not capture this
  - Is cumulative penalty high because of bad learner or bad data?

Notion of Regret

- How do we distinguish between situations where
  - Data is easy and we should expect learner to do very well
  - Data is ridiculous and no learner can do well
- Mistake and Cumulative Penalties do not capture this
  - Is cumulative penalty high because of bad learner or bad data?
- Solution: compare learner against a *benchmark*

Notion of Regret

- How do we distinguish between situations where
  - Data is easy and we should expect learner to do very well
  - Data is ridiculous and no learner can do well
- Mistake and Cumulative Penalties do not capture this
  - Is cumulative penalty high because of bad learner or bad data?
- Solution: compare learner against a *benchmark*
- Common benchmark: set of static predictors

Notion of Regret

- How do we distinguish between situations where
  - Data is easy and we should expect learner to do very well
  - Data is ridiculous and no learner can do well
- Mistake and Cumulative Penalties do not capture this
  - Is cumulative penalty high because of bad learner or bad data?
- Solution: compare learner against a *benchmark*
- Common benchmark: set of static predictors
- Regret of a learning algorithm $\mathcal{A}$

Notion of Regret

- How do we distinguish between situations where
  - Data is easy and we should expect learner to do very well
  - Data is ridiculous and no learner can do well
- Mistake and Cumulative Penalties do not capture this
  - Is cumulative penalty high because of bad learner or bad data?
- Solution: compare learner against a *benchmark*
- Common benchmark: set of static predictors
- Regret of a learning algorithm $\mathcal{A}$

$$\mathfrak{R}_T(\mathcal{A}; \mathcal{W}) = \sum_{t=1}^{T} \ell_t(\mathbf{w}_t) - \min_{\mathbf{w} \in \mathcal{W}} \sum_{t=1}^{T} \ell_t(\mathbf{w})$$

- The algorithm gets to switch predictors, the benchmark gets to see the entire data

Notion of Regret

- Holy grail of online learning: vanishing regret

$$\mathfrak{R}_T(\mathcal{A}) = o(T)$$

- Equivalently

$$\lim_{T\to\infty} \frac{1}{T}\mathfrak{R}_T(\mathcal{A}) = 0$$

$$\lim_{T\to\infty} \left[\frac{1}{T}\sum_{t=1}^{T} \ell_t(\mathbf{w}_t) - \min_{\mathbf{w}\in\mathcal{W}} \frac{1}{T}\sum_{t=1}^{T} \ell_t(\mathbf{w})\right] = 0$$

- Ability to compete with the best predictor in hindsight !!

Up Next

- Brief Introduction to Convex Analysis
- Brief Introduction to Probability Theory

- Online parameter estimation
- Online classification
- Online regression
- Prediction with expert help