When do you use a rules system vs. machine learning?

Apr 9, 2018

Qcon.ai

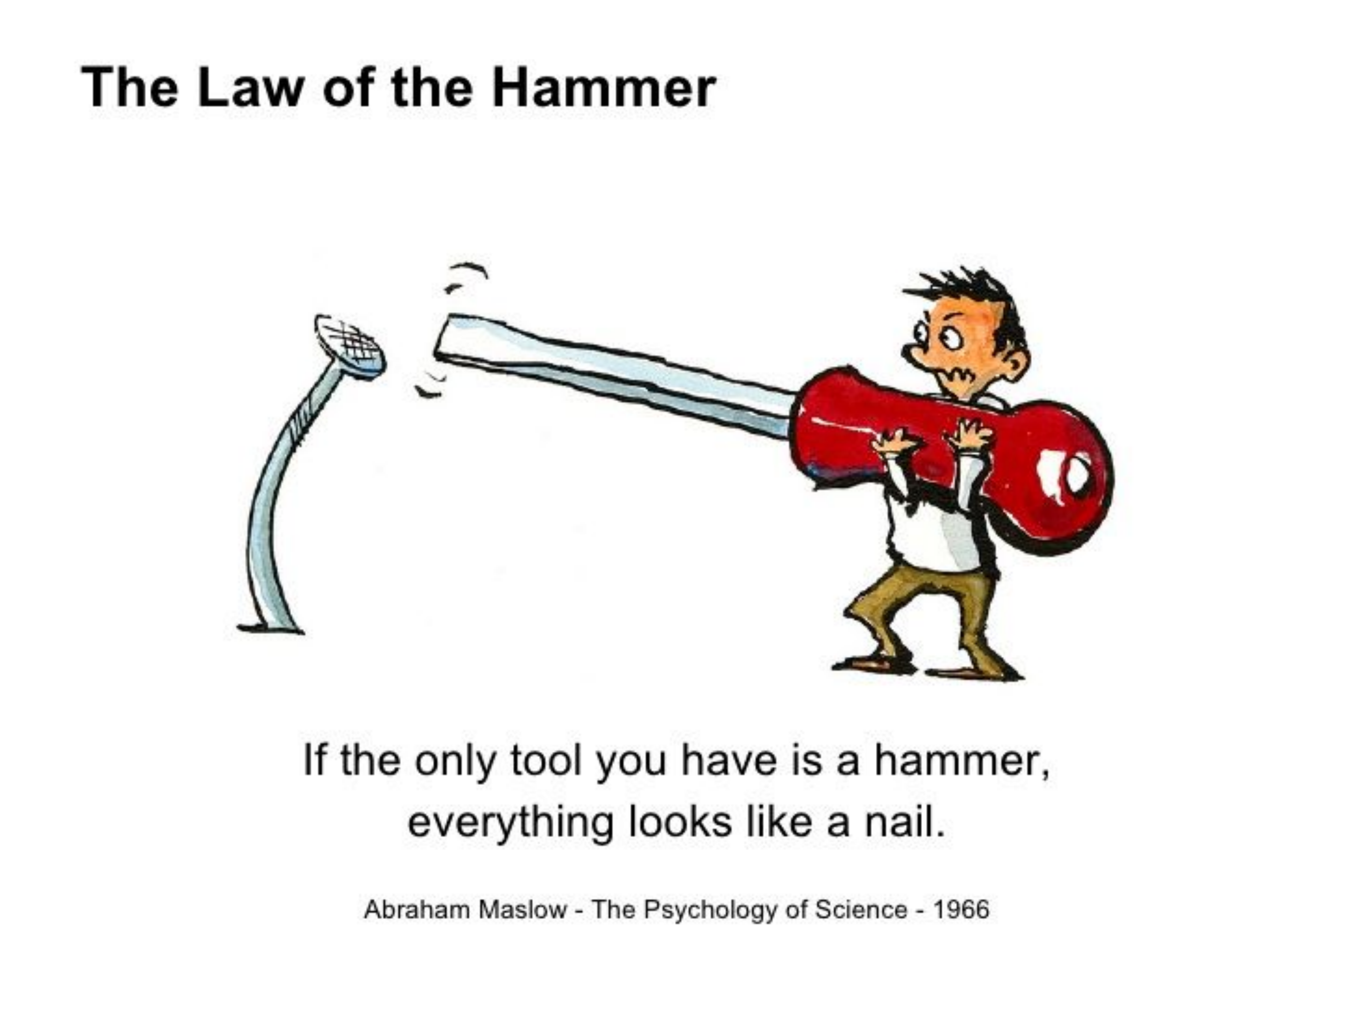

The Law of the Hammer

If the only tool you have is a hammer,
everything looks like a nail.

Abraham Maslow - The Psychology of Science - 1966

When to use what?

Optional sub-title / TL;DR

Criteria	Rules	Machine learning	Deep learning
Interpretability	✓	✓	
Accuracy		✓	✓
Maintenance		✓	
Speed of execution	✓		

Fraud example

What does fraud at Coinbase look like?

Fraudster mashes up multiple identities

1. Fraudster steals Alice’s bank or card info

2. Steals Bob’s Identity

3. Steals Carl’s mobile phone# and verifies it on Coinbase

1. Alice notices fraudulent transaction on her statement

2. Alice disputes the purchase

3. Coinbase returns funds back to Alice

Why does machine learning work in fraud?

We are detecting mismatches across multiple identities, making it harder to create the perfect synthetic identity

Supervised model assigns a risk score to each user and consequently a purchase power

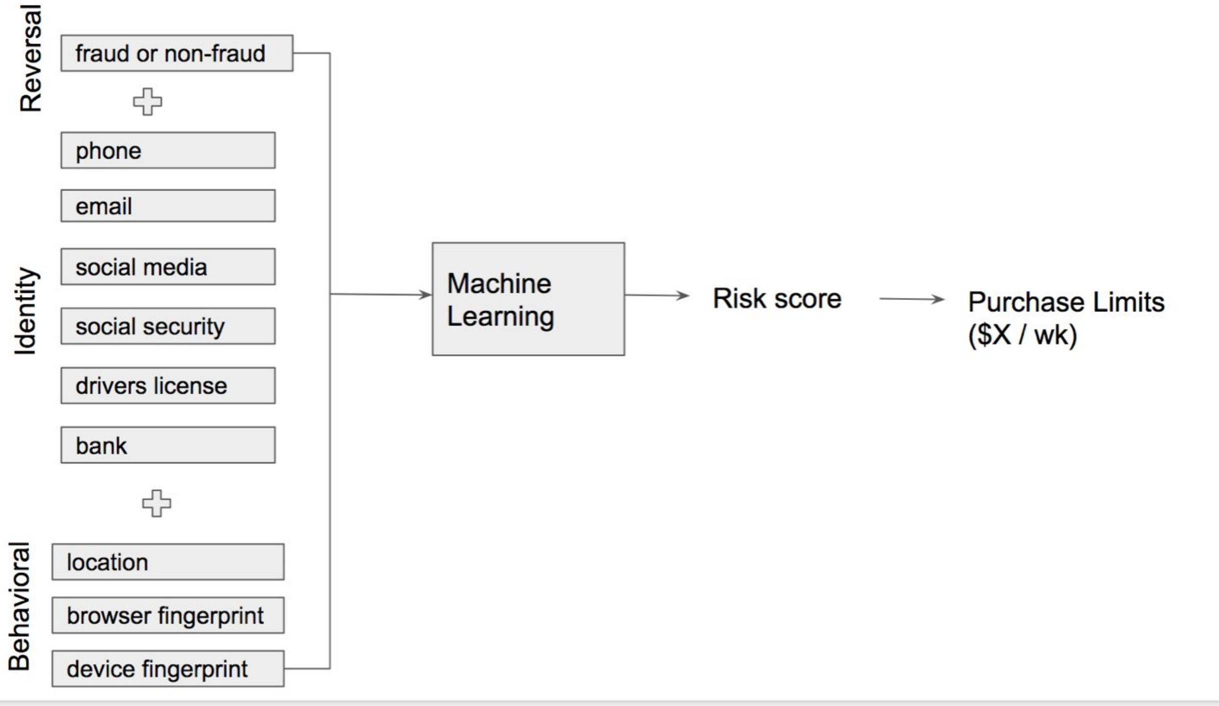

Why does machine learning work in fraud?

Broken window theory ⇒ Track velocity of signals

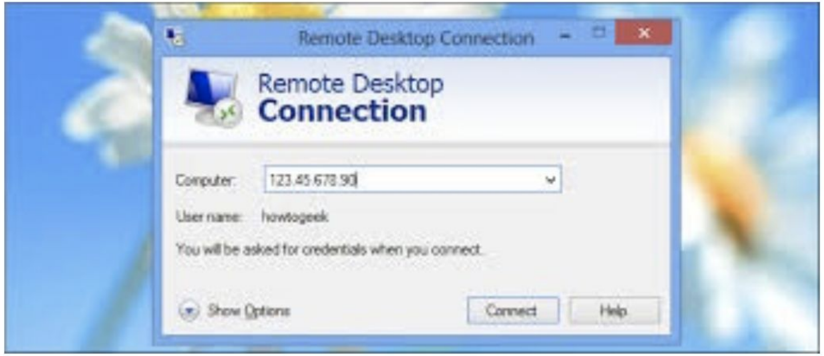

Computer: 123.45.678.90

User name: howtogeek

You will be asked for credentials when you connect.

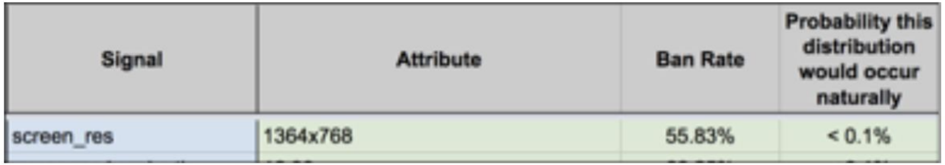

Signal	Attribute	Ban Rate	Probability this distribution would occur naturally
screen_res	1364x768	55.83%	< 0.1%

When does rule based system make sense?

1. Sometimes business use case limits scope of machine learning:
   - Coinbase customers hold their account limits as sacrosanct.
   - This implies we can’t update our machine learning models as frequently as we’d like.
   - However, fraud is dynamically evolving. How do you catch up?

2. Lack of adequate training data

When does rule based system make sense - part 1

Business use case limits scope of machine learning

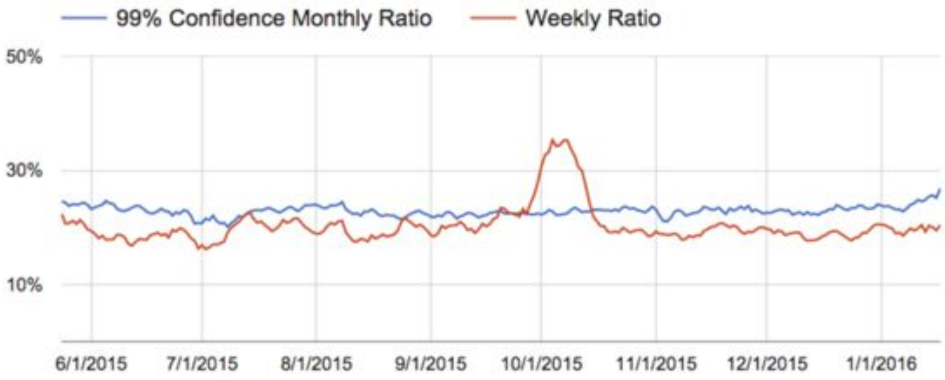

1. Identify patterns

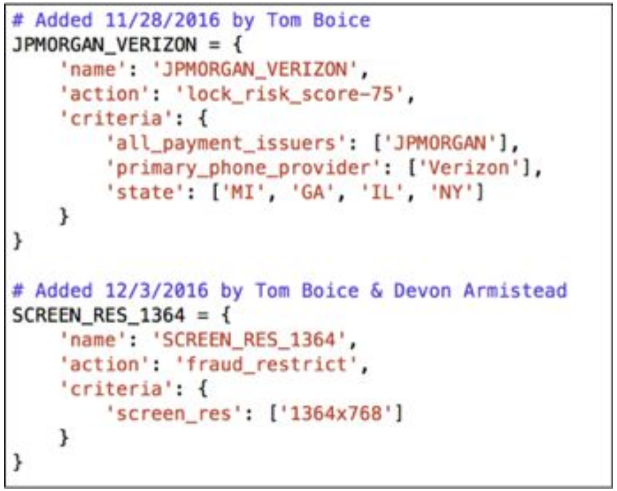

```
# Added 11/28/2016 by Tom Boice
JPMORGAN_VERIZON = {
    'name': 'JPMORGAN_VERIZON',
    'action': 'lock_risk_score-75',
    'criteria': {
        'all_payment_issuers': ['JPMORGAN'],
        'primary_phone_provider': ['Verizon'],
        'state': ['MI', 'GA', 'IL', 'NY']
    }
}

# Added 12/3/2016 by Tom Boice & Devon Armistead
SCREEN_RES_1364 = {
    'name': 'SCREEN_RES_1364',
    'action': 'fraud_restrict',
    'criteria': {
        'screen_res': ['1364x768']
    }
}
```

2. Apply restrictions

Things to watch out for in rules systems

Watch out for technical debt.

- Monitor impact of each rule
- Retire rules as quickly as you add them

When does rule based system make sense - part 2

Lack of representative training data implies machine learning won’t work.

Account takeovers are very infrequent compared to fraud

- Only 100s of examples, each of which exploits a unique vulnerability
- This means you can’t use machine learning to extrapolate and generalize

1. Detect suspicious login or withdrawal patterns

User behavior data is the key element here

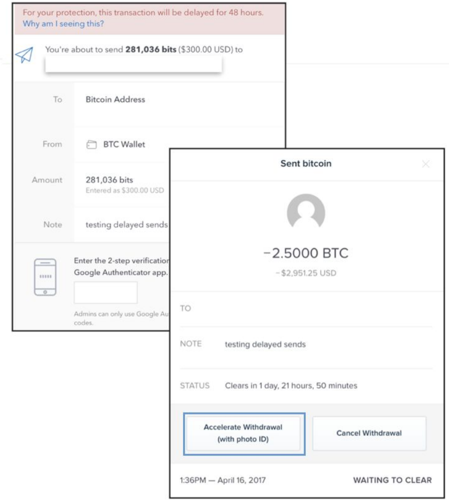

For your protection, this transaction will be delayed for 48 hours.
Why am I seeing this?

You're about to send **281,036 bits** ($300.00 USD) to

To Bitcoin Address

From BTC Wallet

Amount 281,036 bits
Entered as $300.00 USD

Note testing delayed sends

Enter the 2-step verificatio
Google Authenticator app.

Admins can only use Google Au
codes.

Sent bitcoin

−2.5000 BTC
−$2,951.25 USD

TO

NOTE testing delayed sends

STATUS Clears in 1 day, 21 hours, 50 minutes

1:36PM — April 16, 2017 WAITING TO CLEAR

2. Delay suspicious withdrawals

3. Allow false positives to accelerate a delayed withdrawal via further proof of authenticity

Summary

TL;DR — Use the right tool for the right job

- Depending on your use case, rules systems may be quicker to build and iterate
- It is important to know when to transition from rules to machine learning:
  - Too many rules
  - People are afraid to remove rules