rapidminer

The Model Impact Epidemic

Less than 1% of models have their desired impact

Worldwide spending on AI is expected to grow to $37.5 billion dollars this year, however there's been increasing awareness of a set of issues that are preventing many organizations from operationalizing the models they're building and maximizing the investments that are being made. There's an epidemic that's preventing models from making it into deployment where they can actually have an impact.

RapidMiner has analyzed the issue and documented our findings in an attempt to understand the root causes, so we know how to address them to stop the epidemic in its tracks. It's important to begin by quantifying how many models are actually making it into production across all enterprises so we can identify the biggest bottlenecks.

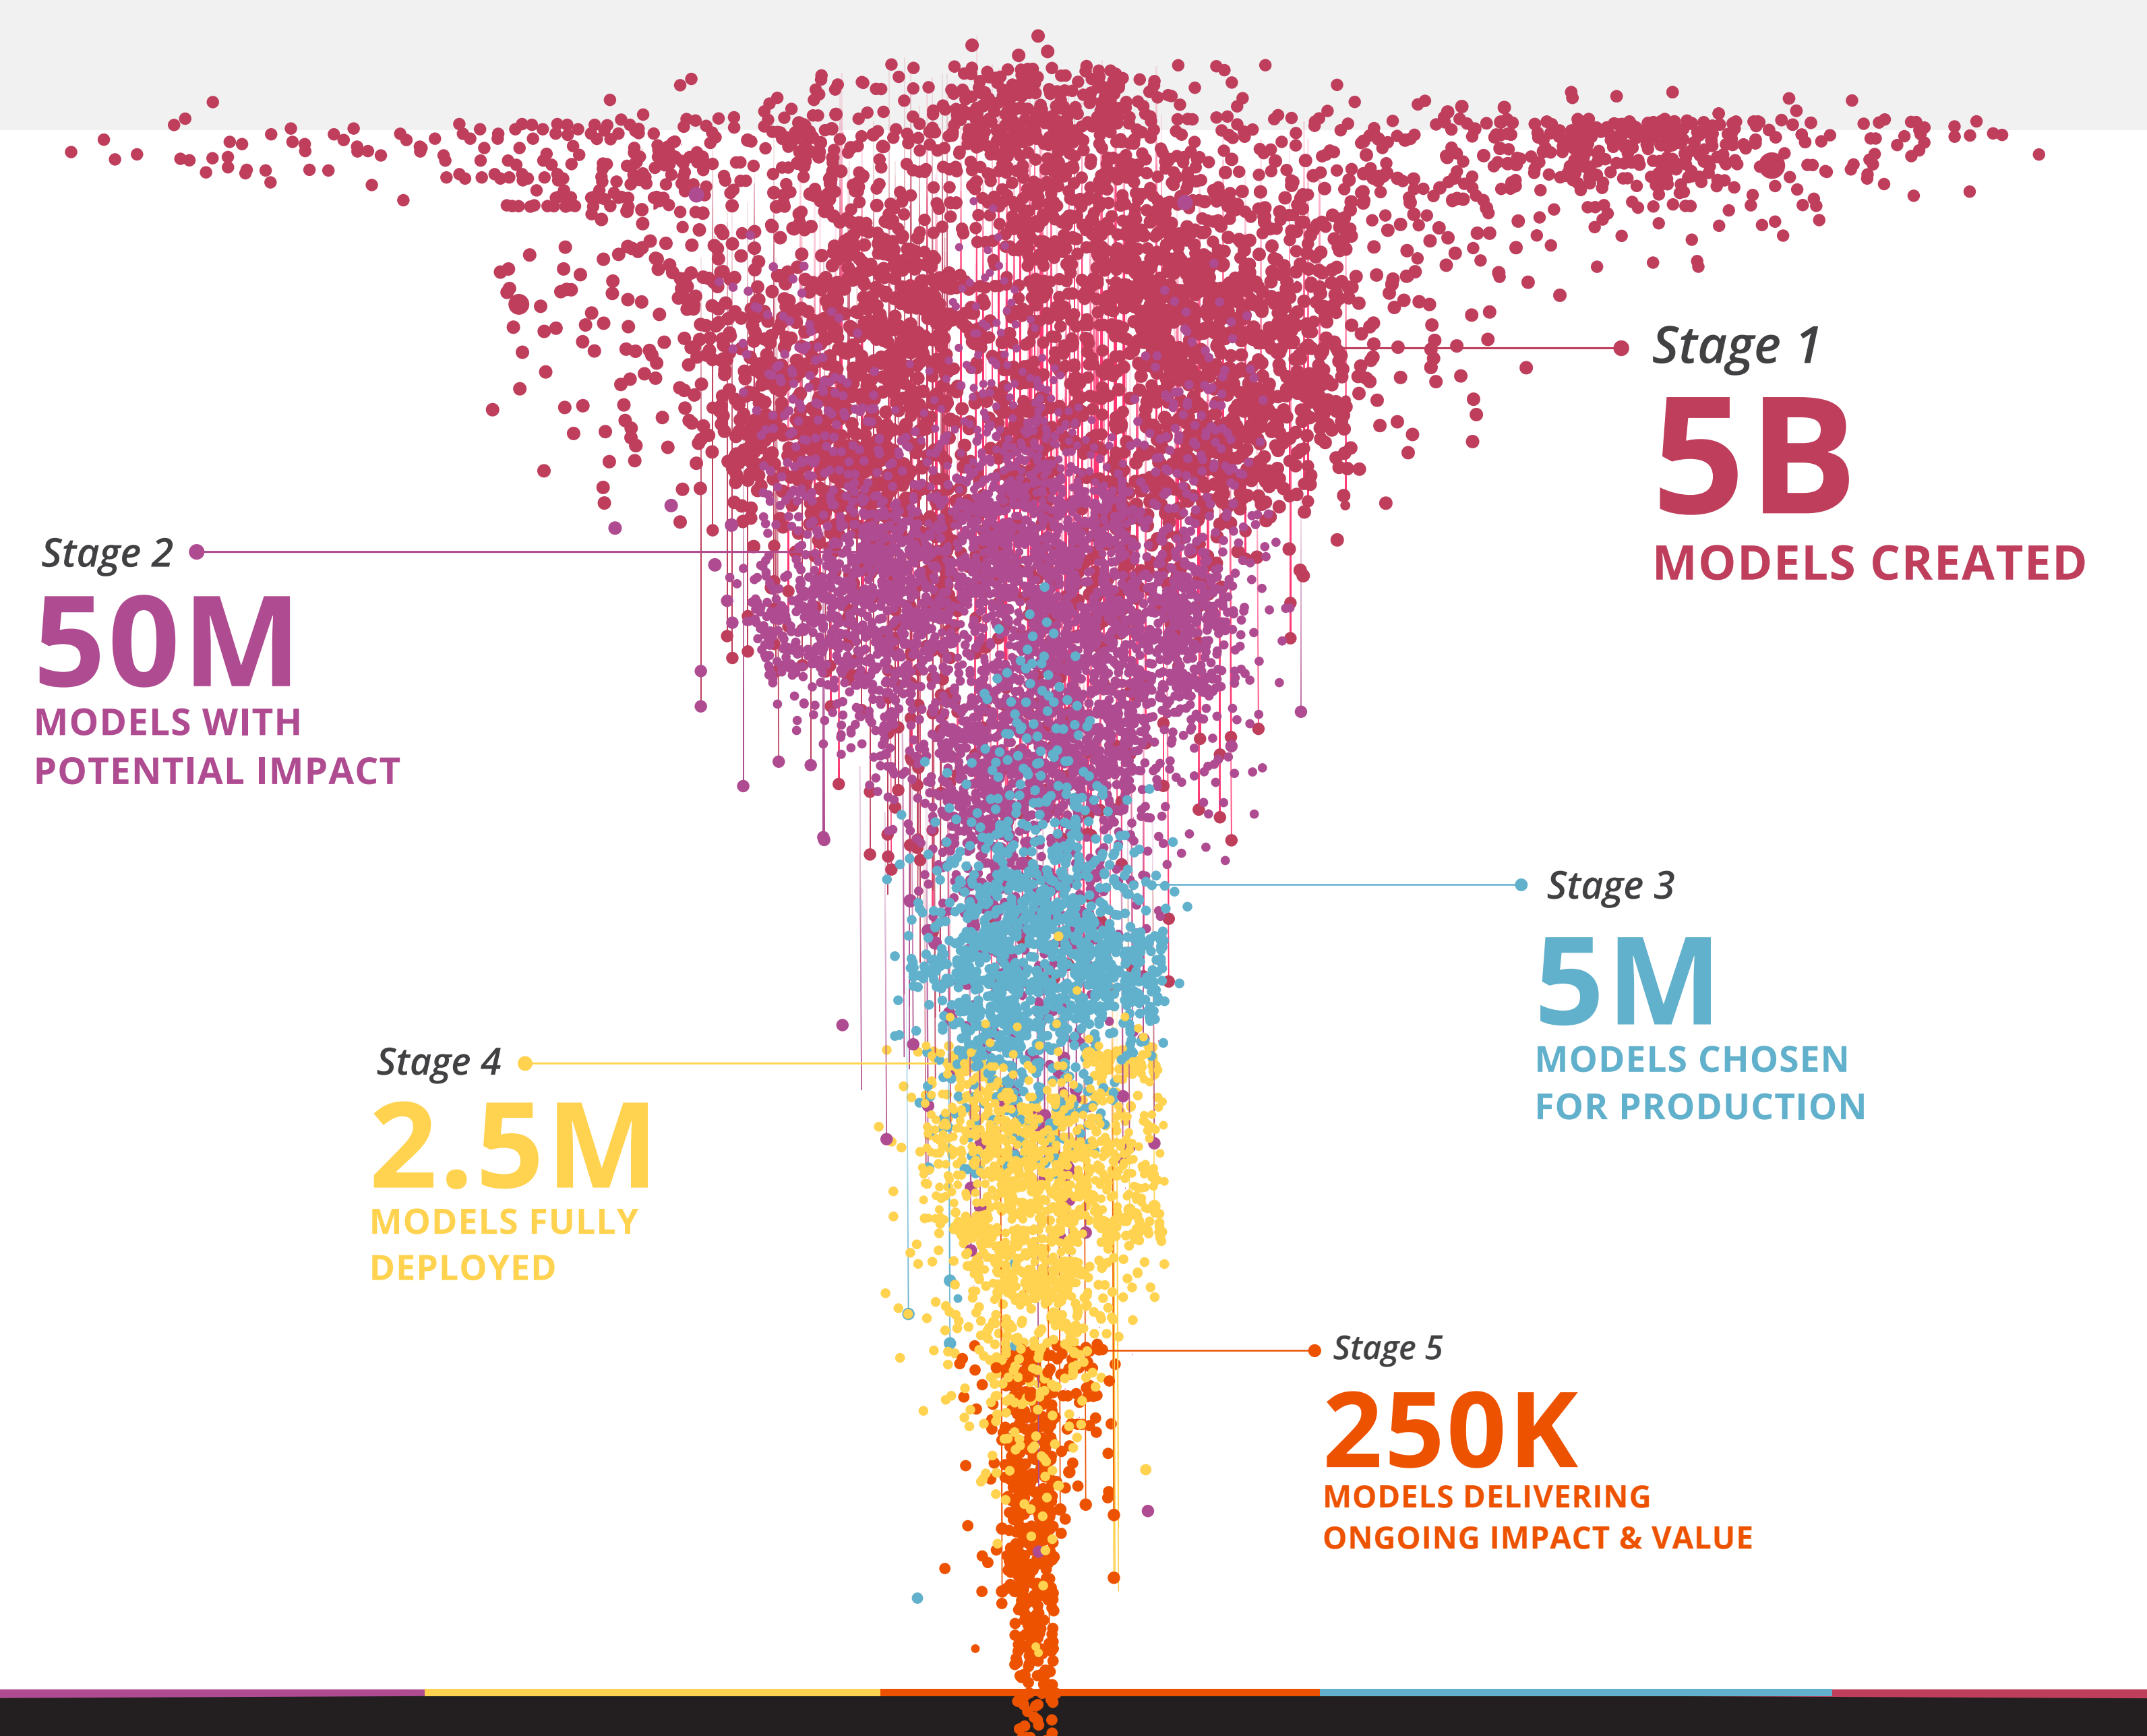

The 5 Stages of the Epidemic

5,000,000,000

Stage 1

CREATING MODELS

This is the starting point for every project, it requires you to prototype many models to find the right fit for your use case. Usually hundreds of models are created and evaluated during this stage per project.

$$= \underbrace{10^6}_{\text{1M data scientists}} \times \underbrace{5}_{\text{use cases per project per year}} \times \underbrace{1000}_{\text{model candidates per use case}}$$

50,000,000

Stage 2

FINDING MODELS WITH IMPACT

This is where you must evaluate, validate and test the hundreds of models to find the one that can actually deliver enhanced decision-making to drive cost reduction, revenue gain, or risk reduction.

$$= \frac{5 \times 10^9 \text{ model candidates}}{1000 \text{ model candidates per use case}}$$

5,000,000

Stage 3

CHOOSING MODELS FOR PRODUCTION

This is where the potential impact of the model is shared and discussed. As buy-in is garnered across stakeholders, resources are committed to deploy a model into production and timelines are put in place. Will this be an interactive dashboard? A weekly report? A part of a website? Estimates state that only 1 out of every 10 data science projects are chosen for production.

$$= \underbrace{50 \times 10^6}_{\text{Models with impact}} \times \underbrace{10\%}_{\text{Models chosen for production}}$$

2,500,000

Stage 4

DEPLOYING MODELS

This is where the rubber hits the road and IT, DevOps and domain experts work to get the model fully operational within a business workflow. This often requires a diverse set of experience, including coding, computing resources management, and a deep understanding of the business workflow that's being enhanced. Estimates state that less than half of data science projects that are supposed to be in production end up fully deployed.

$$= \underbrace{5 \times 10^6}_{\text{Models for production}} \times \underbrace{50\%}_{\text{Models fully deployed}}$$

250,000

Stage 5

DELIVERING ONGOING IMPACT AND VALUE

Once models are deployed the hard work is over, but the journey is not. It's still critical to monitor models for performance, business impact and potential drift issues. It may be necessary to change the actual model if alarming shifts occur. Continuous monitoring requires additional resources that many business don't account for, but is necessary to deliver ongoing impact.

$$= \underbrace{2.5 \times 10^6}_{\text{Models in production}} \times \underbrace{10\%}_{\text{Models delivering ongoing value}}$$

Ultimately, less than 1% of models that can have an impact on a business are actually deployed & delivering value.

The Epidemic is Making Headlines

The numbers may vary slightly depending on who you ask, but the story is the same… and it is troubling.

$37B *Expected worldwide AI spend this year*
- IDC

47% *More than half of data science projects are not fully deployed*
- Gartner

13% *of data science projects, or just one out of ten, actually make it to production*
– Venture Beat

Root Causes of the Epidemic by Stage

Stage 1
Creating Models
- You must try 100s of models for every new project (No free lunch theorem)
- Auto ML reduces human effort, but still…
- …this many models is hugely impacting energy expenditure
- If we use these resources we better make it count

Stage 2
Finding Models with Impact
- The right data doesn't exist or can't be accessed
- Rookie mistakes can be made by 'citizen data scientists'
- Models created without domain expertise may not make sense
- Data scientists can overinvest into solution paths and overlook simpler approaches (local optimum)
- Conformity bias – deep learning may not always be the best solution path

Stage 3
Choosing Models for Production
- Black box models
- Explaining models is hard and not always part of a data scientists core skillset
- May require separate analysis to show business impact, which is needed for buy-in
- Concerns over having the resources to deploy, monitor, and manage models
- Some decision-makers refuse to believe that models can trump intuition

Stage 4
Deploying Models
- You have to connect to other pieces of IT infrastructure, and you may not have permission
- You may need to rely on other people, and it may not be their priority
- Change requests mean you need to continue the perpetual convincing
- Creating proper solutions for model monitoring and management can be more work than building the model itself

Stage 5
Delivering Ongoing Impact and Value
- People don't understand the model or its predictions
- Need to do more convincing and explaining
- No ongoing monitoring of drift and bias
- Wrong validation in earlier stages may lead to less or even negative business impact
- Problems go undetected due to lack of model operations and management

The Spiral of Disillusionment

Some of the most widespread and avoidable issues that prevent good AI models from delivering the desired results occur between stages 4 & 5. This is because most organizations don't have sufficient model operations resources in place and it leads to bad experiences. We call this the 'Spiral of Disillusionment.'

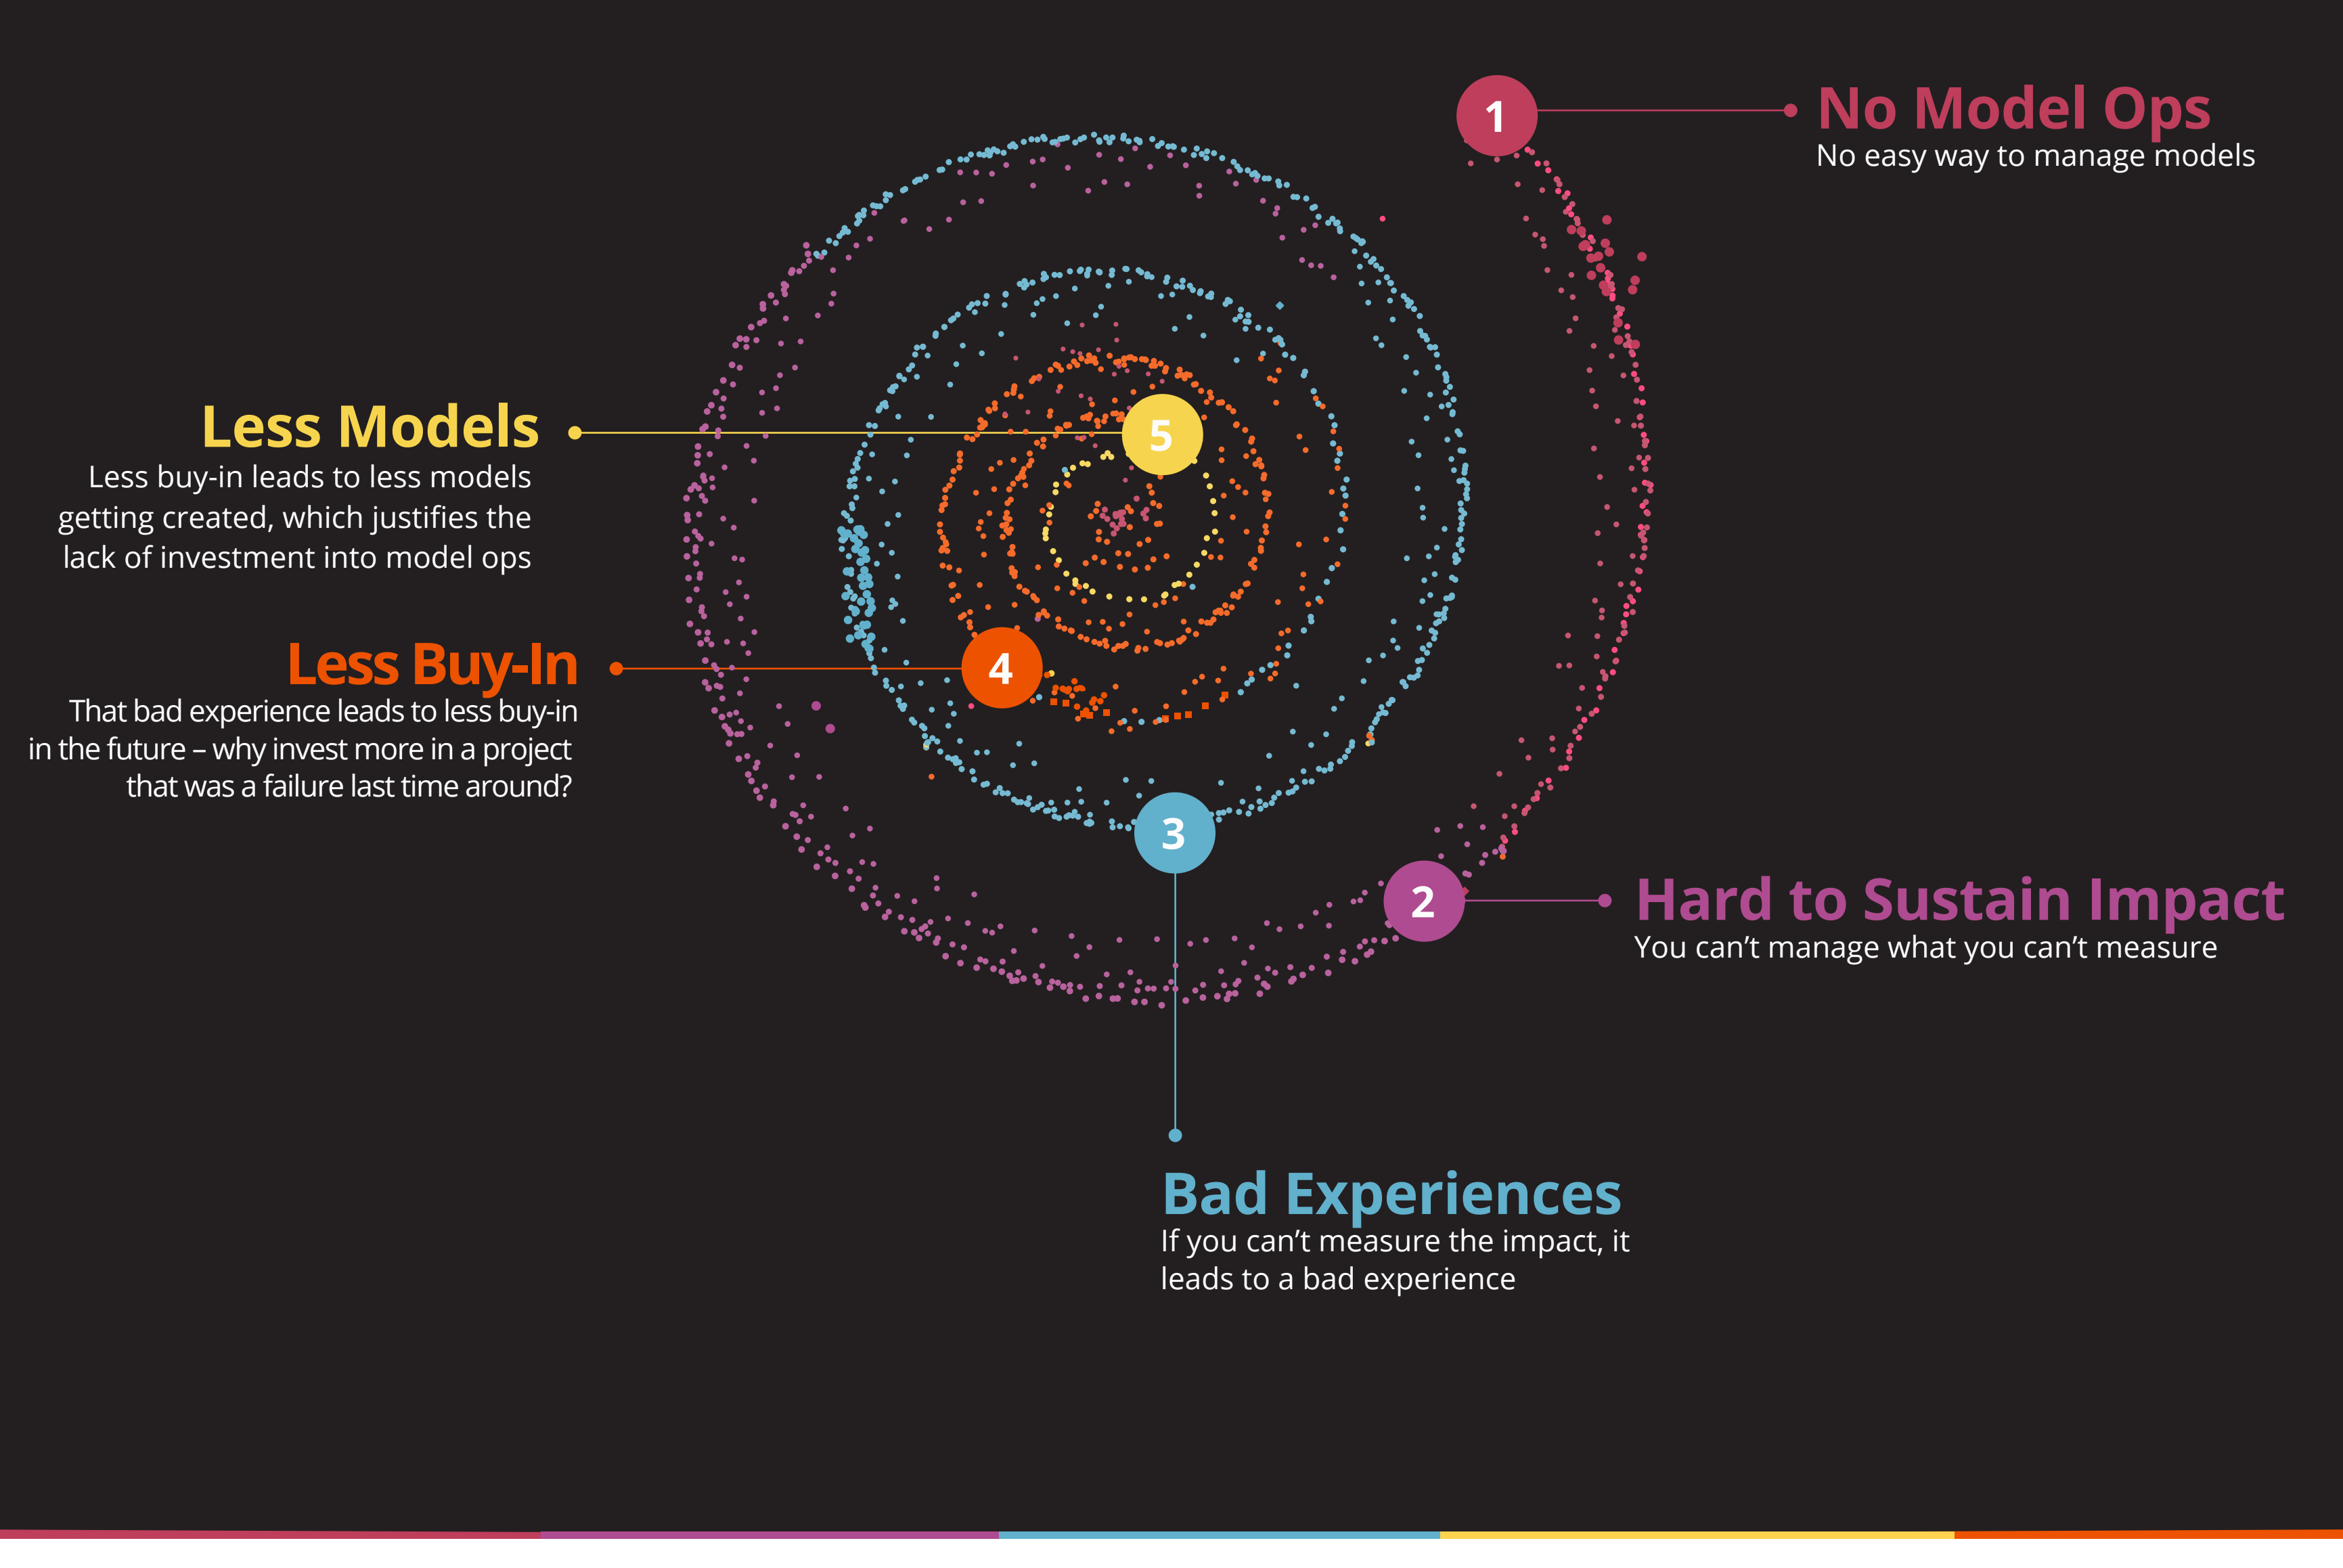

No Model Ops
No easy way to manage models

Hard to Sustain Impact
You can't manage what you can't measure

Bad Experiences
If you can't measure the impact, it leads to a bad experience

Less Buy-In
That bad experience leads to less buy-in in the future – why invest more in a project that was a failure last time around?

Less Models
Less buy-in leads to less models getting created, which justifies the lack of investment into model ops

How RapidMiner 9.4 Cures the Epidemic

Model Ops

An easy way for business users to put models into production. Users can automatically create robust scoring processes, integrate with other IT systems, manage and monitor performance on a model leaderboard, see and prevent concept drift and bias.

- Reduces skill requirements to deploy worthy models
- Reduces overreliance on IT and DevOps resources to deploy models into production
- Reduces risks of deploying models by easily managing concept drift , bias and streamlining governance
- Closes the feedback loop - create prediction, act, see impact so you can improve models over time

Managed Offerings in the RapidMiner AI Cloud

Allows users to deploy models into production without acquiring and managing infrastructure.

- Elastic cloud-based data science resources on-demand to run complex scoring without massive on-prem server resources

Auto Model Web

A new browser-based version of the proprietary RapidMiner Auto Model technology, built for business users who know their data and use case, but don't have advanced data science background.

- Increases the number of people who can produce accurate models and evaluate them
- Democratizes data science to ensure good initial model quality at volume and scale

Profit-Sensitive Scoring

A unique capability which allows business users to input cost and revenue variables in order for the model to self-optimize for profitability.

- Shows impact of a model so that you can facilitate better business buy-in upfront to get models into production
- Helps to make more confident decisions about whether a model should be in production in the first place

Enhanced Visualizations & New Charts

Helps tell a compelling and intuitive story about data and models.

- Easily explore and explain complex data sets and models to support stakeholder buy-in

rapidminer

RapidMiner is reinventing enterprise AI so that anyone has the power to positively shape the future. We're doing this by enabling 'data loving' people of all skill levels, across the enterprise, to rapidly create and operate AI solutions to drive immediate business impact. We offer an end-to-end platform that unifies data prep, machine learning, and model operations with a user experience that provides depth for data scientists and simplifies complex tasks for everyone else. Our Center of Excellence methodology and the RapidMiner Academy ensures customers are successful, no matter their experience or resource levels. More than 30,000 organizations in over 150 countries rely on RapidMiner to increase revenue, cut costs, and reduce risk.

rapidminer.com

IDC, Worldwide Spending on Artificial Intelligence Systems Will Be Nearly $98 Billion in 2023, 04 Sep 2019.
Gartner, How to Operationalize Machine Learning and Data Science Projects, Erick Brethenoux, Shubhangi Vashisth, et al., 3 July 2018.
Venture Beat, Why do 87% of data science projects never make it into production? 19 July 2019