comet

REPORT

2021 Machine Learning Practitioner Survey

Too Much Friction, Too Little ML

TABLE OF CONTENTS

Executive Summary: Too Much Friction, Too Little ML

Artificial intelligence (AI) is encountering yet another hurdle to delivering value, in the form of friction among and between teams. As machine learning (ML) has delivered outsized business value on common use cases such as detecting fraud, recommending products, and predicting customer churn, more companies are seeking to apply it to innovative use cases – only to find the tools and processes to be disconnected, unreliable, and unscalable.

Challenges related to people, process, and tools are creating friction that makes it difficult to track model training runs and results, collaborate with colleagues, and iterate faster during the complex process of developing machine learning. This friction can cause delays in ML development that delay or halt model deployment to production.

In a survey of 508 machine learning practitioners that included data scientists and engineers:

- **58% reported they are using manual tools to track ML experiments**
- **68% admitted abandoning 40-80% of their experiments in the past year**
- **47% reported their teams require 4-6 months to deploy a single ML project**
- **88% of respondents' companies have an annual budget of less than $75,000 for machine learning tools and infrastructure**
- **63% said their organizations would increase ML budgets for 2022**

But there's good news; pioneering enterprise companies have developed ML development frameworks and processes, and many of these are available via open source.

ML practitioners should look for tools that:

1. Support integration and customizability
2. Scale to handle the intense demands of production environments
3. Apply across the complete lifecycle of ML development, from data preparation and model training to deployment and monitoring models in production

Survey Methodology

Market research company Censuswide conducted a one-time survey of 508 U.S. machine learning practitioners across industries. Respondents included data scientists and engineers who agreed to be surveyed.

The survey was administered online between September 13 and October 4, 2021. Respondents answered multiple-choice questions about ML development and factors affecting their teams' ability to deliver the level of business value their companies expect from data-driven initiatives.

Censuswide abides by and employs members of the Market Research Society, which is based on ESOMAR principles. ESOMAR is the global voice of the data, research and insights community, representing individual professionals and companies who provide or commission data analytics and research.

Survey Questions and Results

1. What industry is your company in?

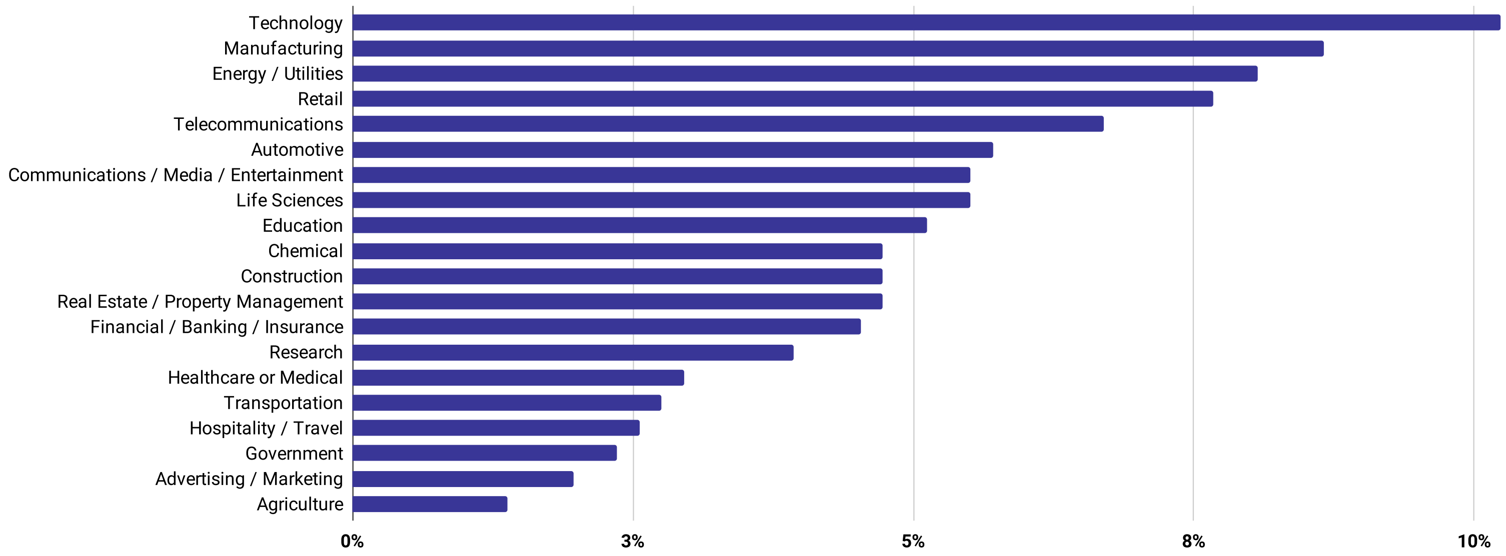

2. Approximately what was your company's revenue in 2020?

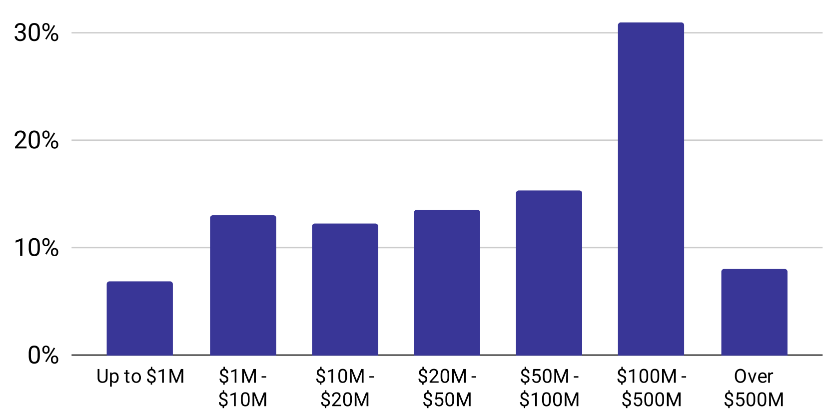

3. To what extent did your company experience revenue growth or decline in the last 12 months?

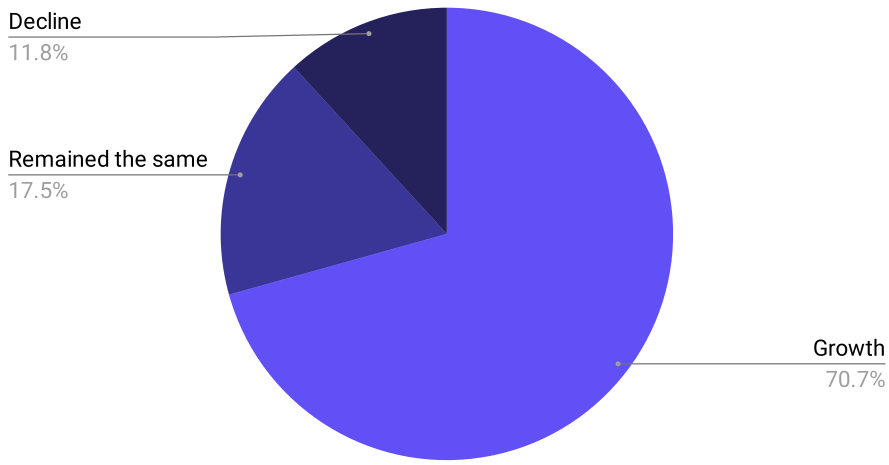

4. How many teams within your company work on your machine learning projects?

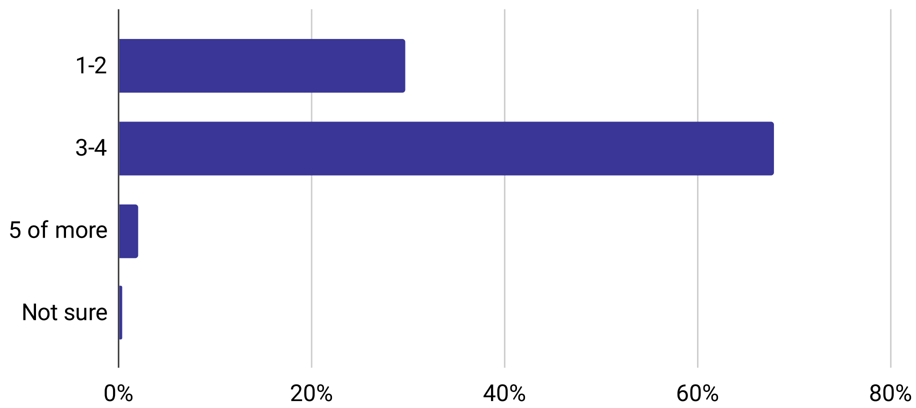

5. What tools, if any, does your team use for tracking the ML development process?

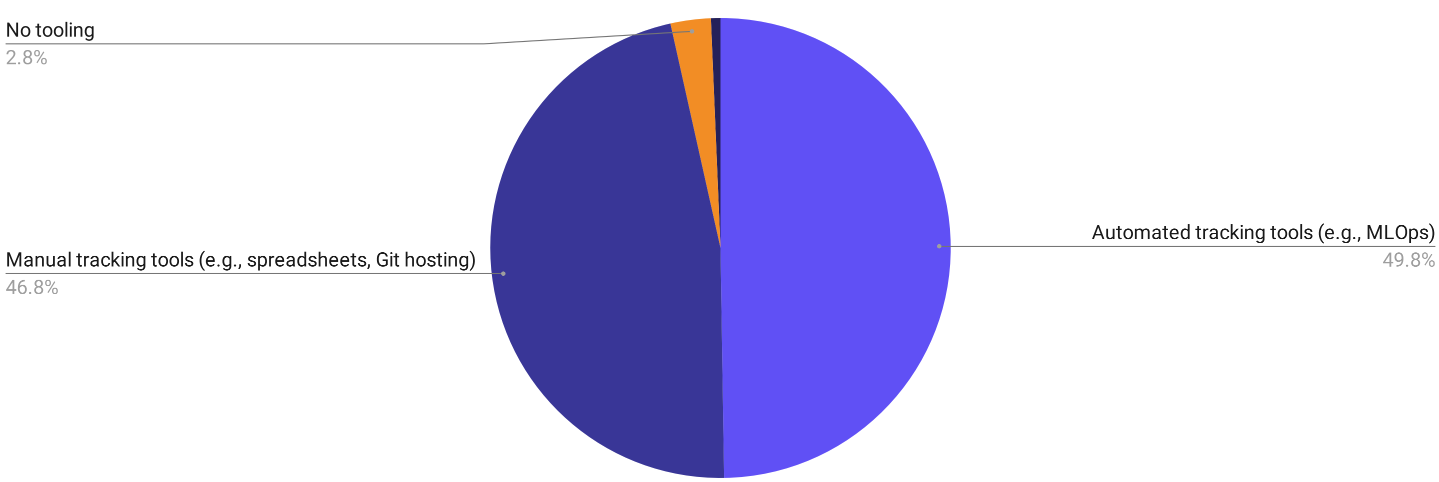

6. Roughly how long does it take your team to deploy a single machine learning project?

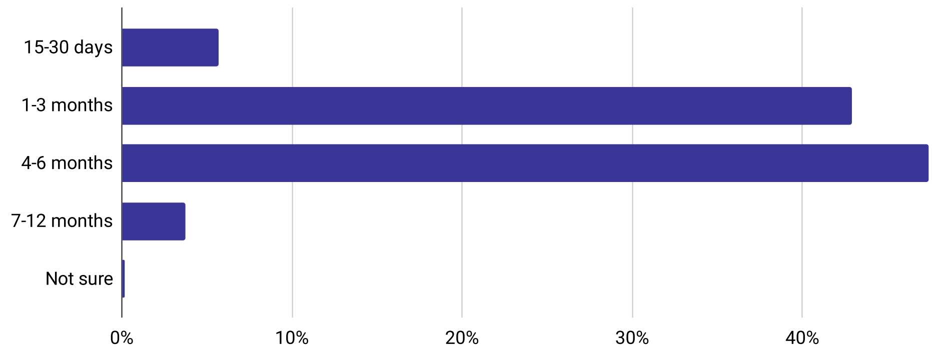

7. Roughly how many of your machine learning experiments had to be scrapped due to mismanagement in the past year?

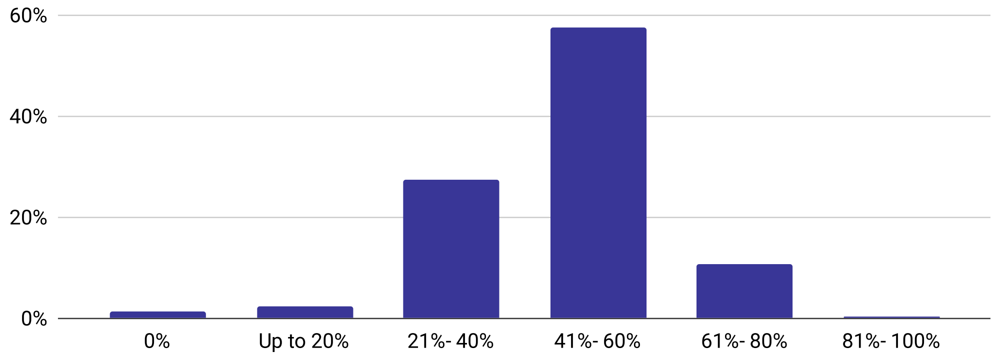

7a. Why did your ML experiments have to be scrapped? (select all that apply)

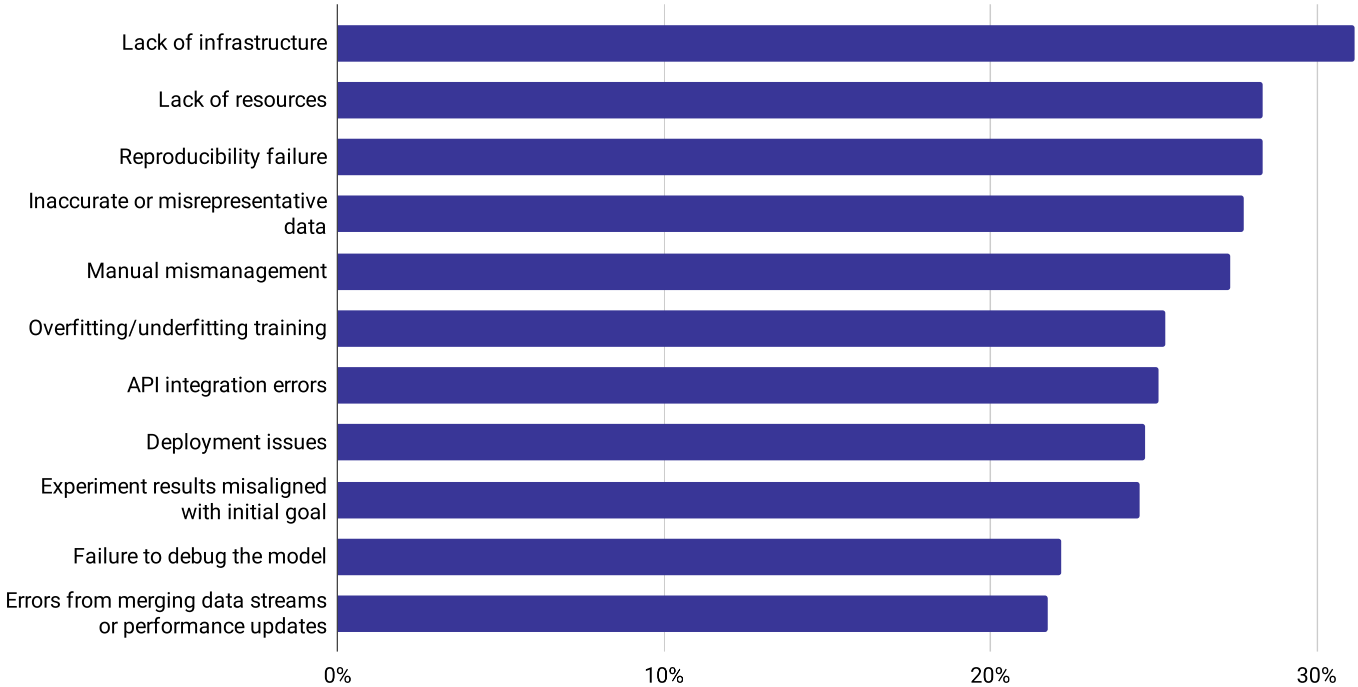

8. How does your company store data used in machine learning experiments?

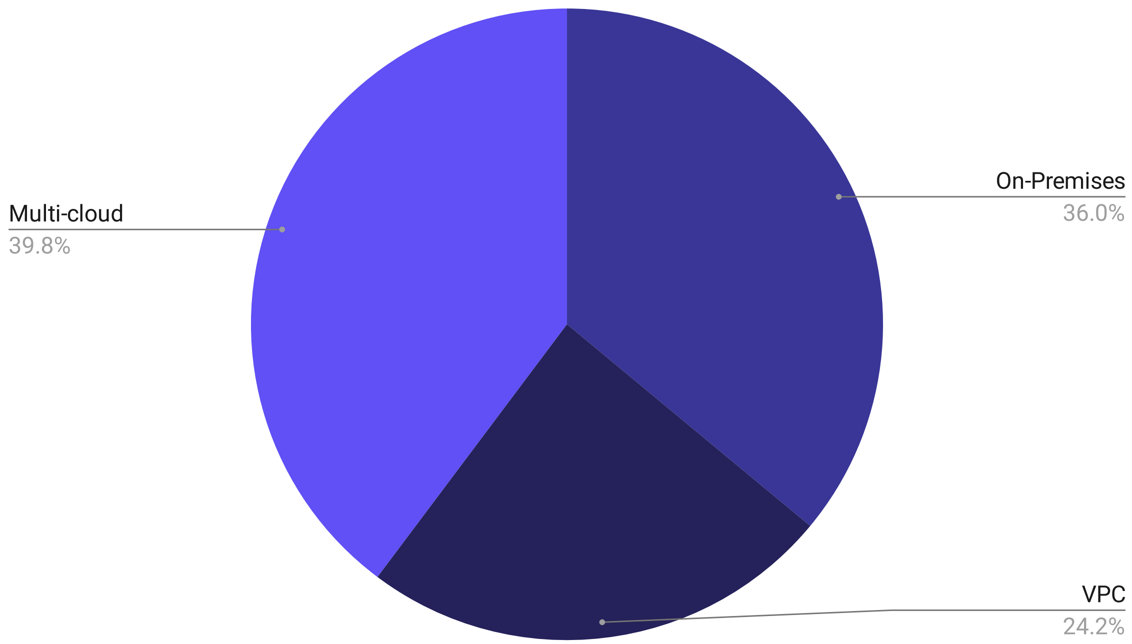

9. What is the biggest pain point slowing down your AI roadmap?

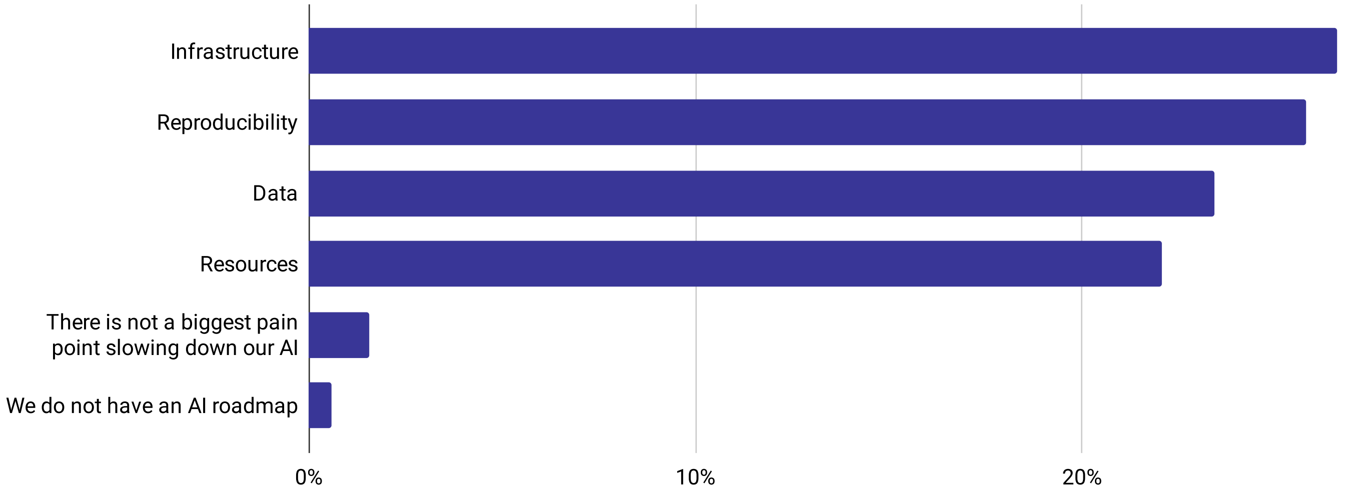

10. What does your company currently use machine learning projects for? (select all that apply)

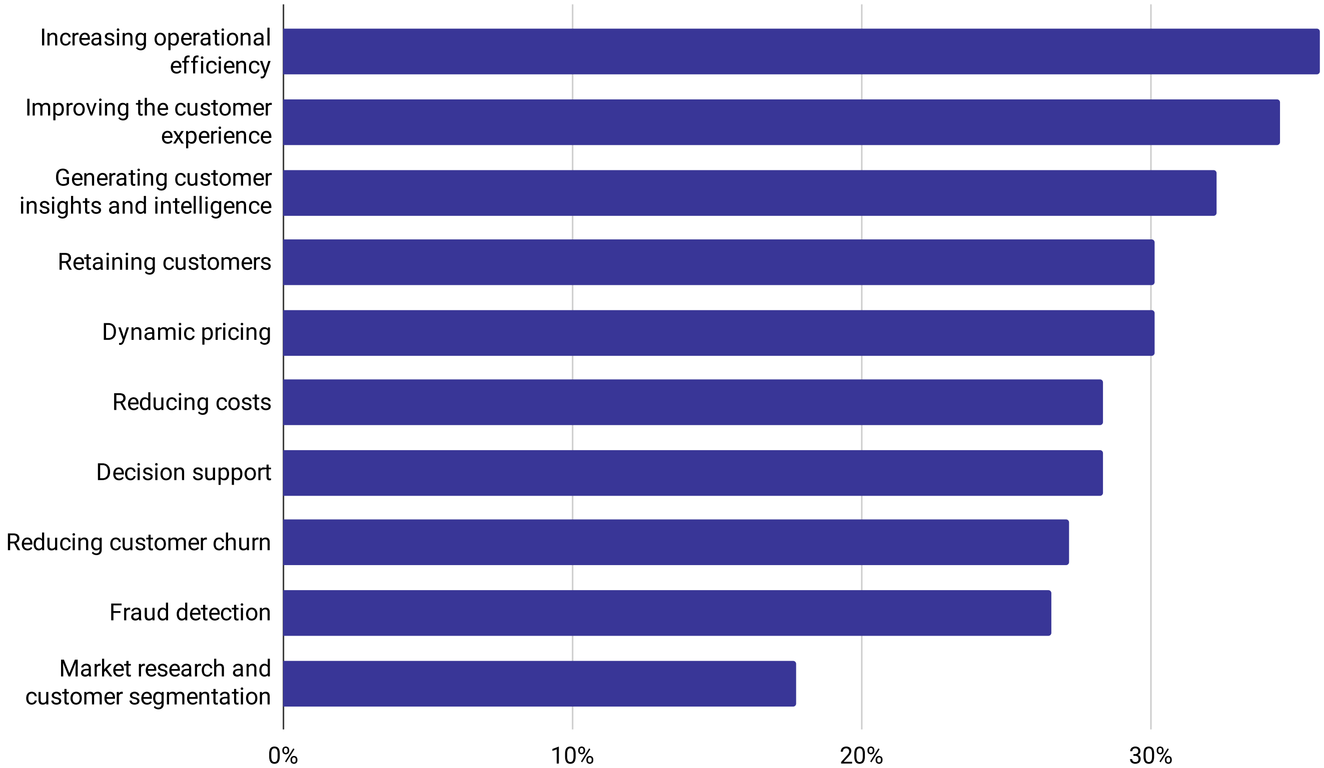

11. In what ways would you like your company to use machine learning (that it is currently not) in the next 1-3 years? (select all that apply)

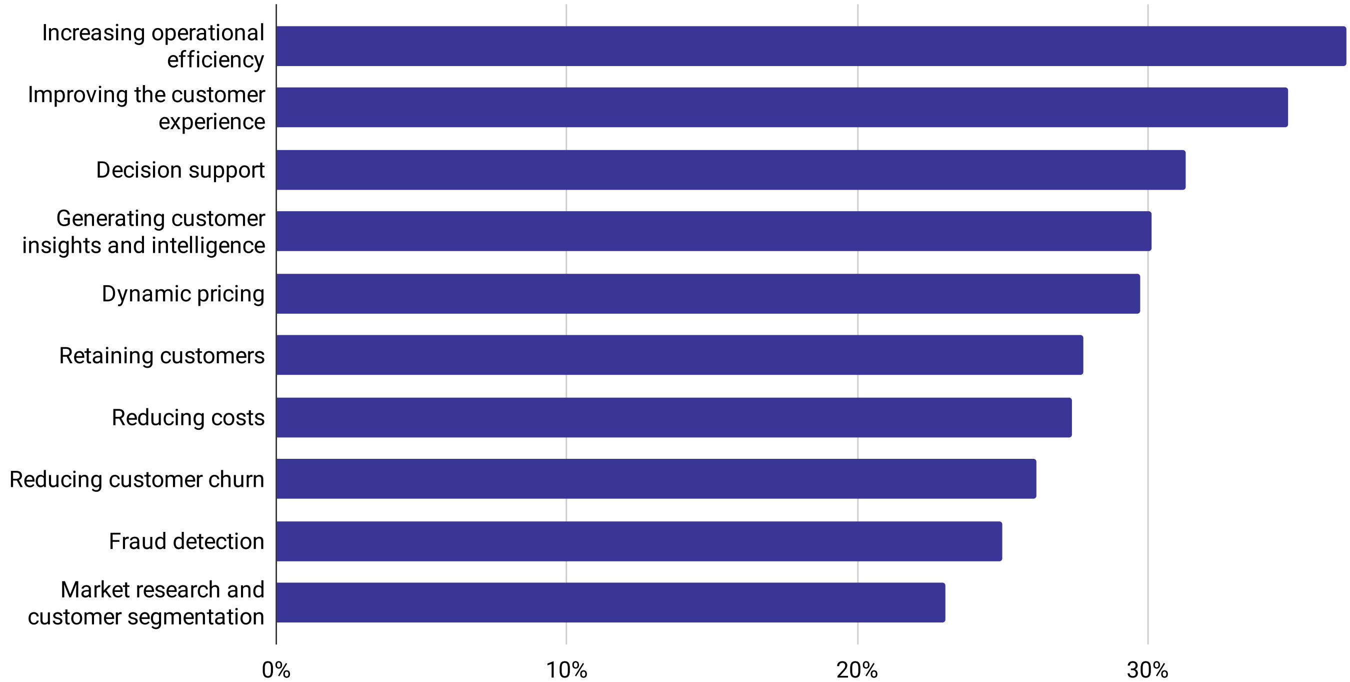

12. Approximately how many machine learning projects does your company have in production today?

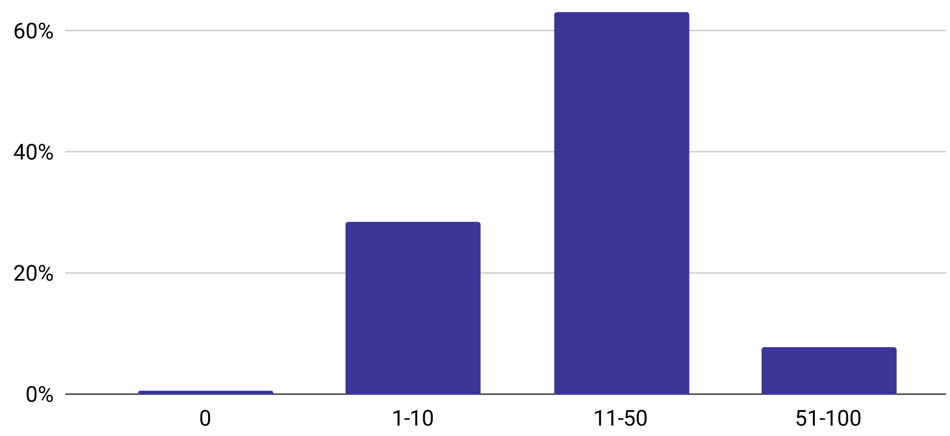

13. How many machine learning projects do you want to have in production in 6-12 months?

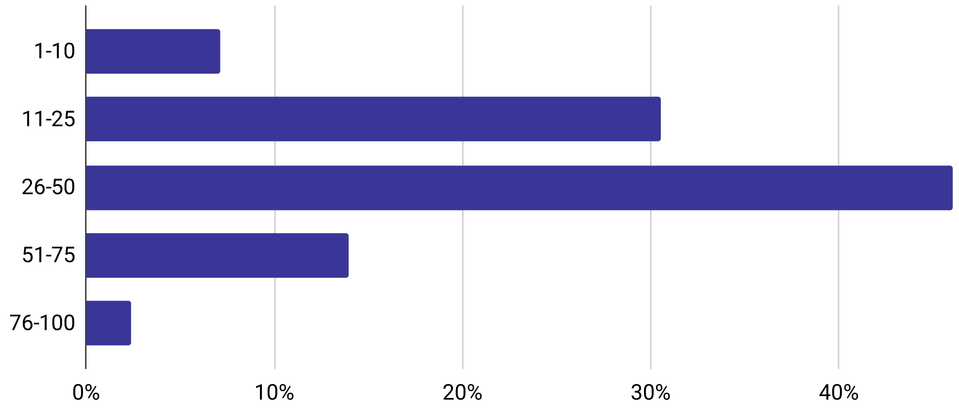

14. What is your company's current annual budget for machine learning?

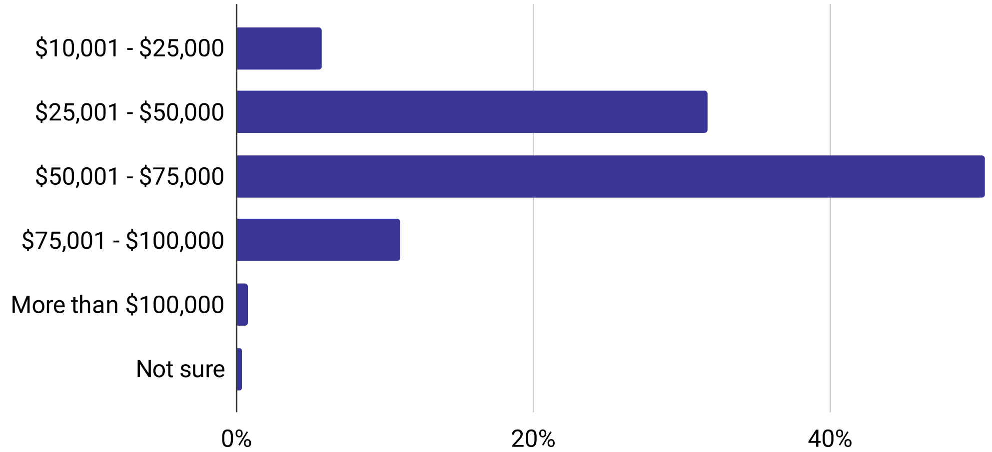

15. How will your company's machine learning budget change in the next 12 months?experiments?

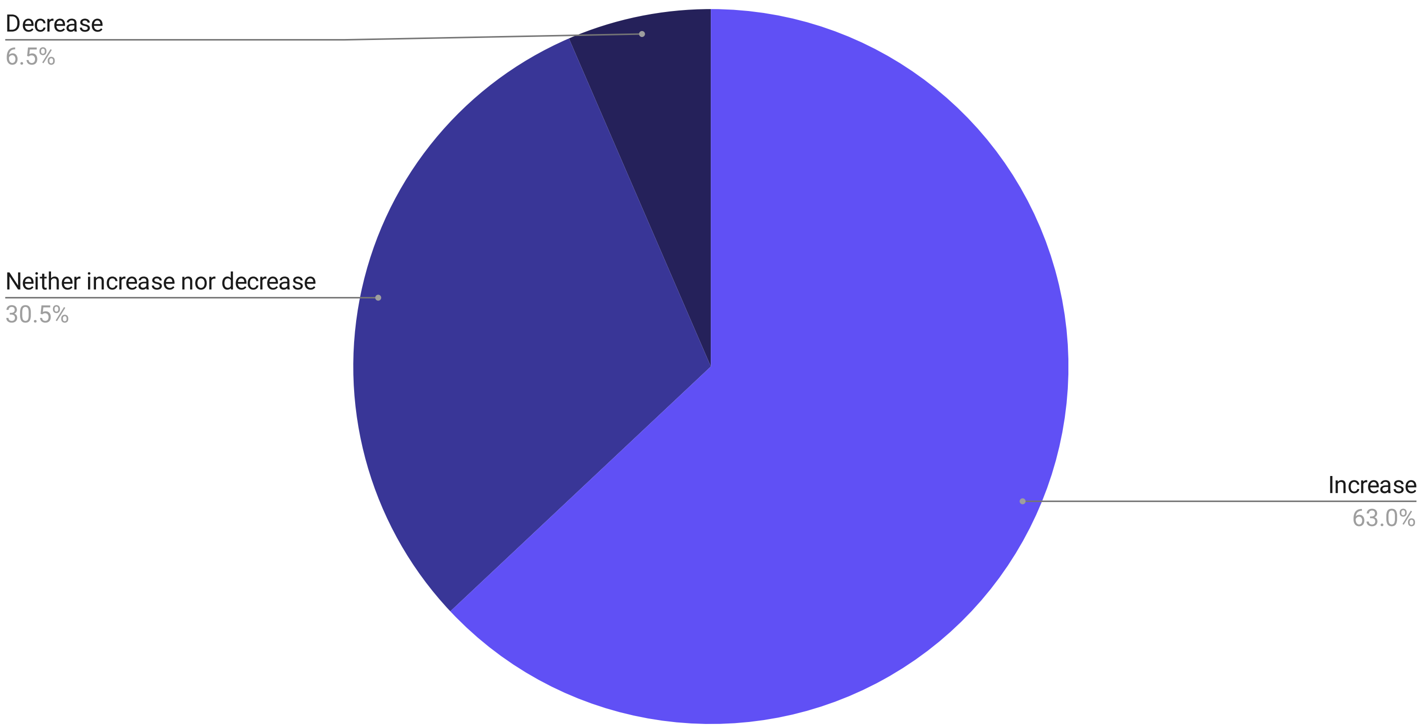

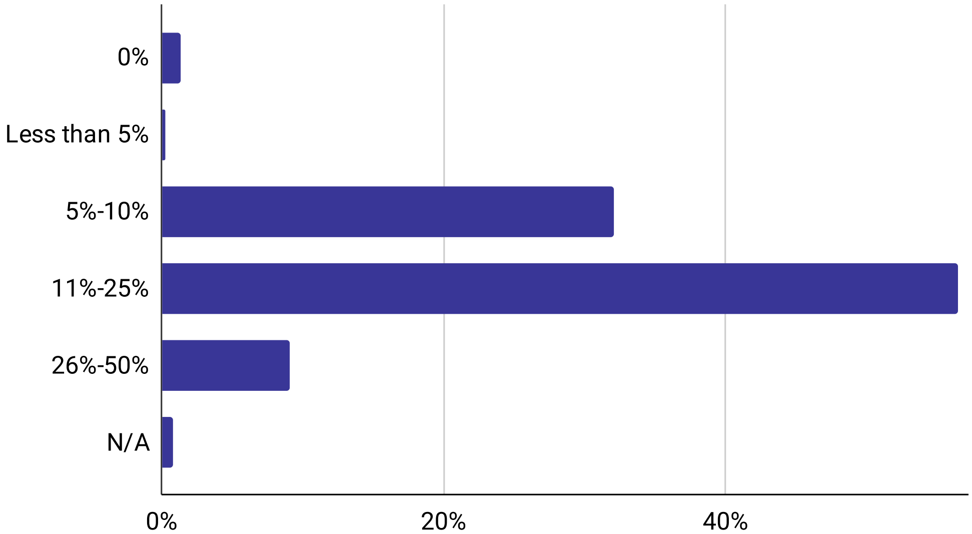

16. How many of your ML projects failed once they were launched from an experiment into real-world production in the last year?

Key Survey Findings and Analysis

According to survey respondents, there are four challenges companies face in developing machine learning:

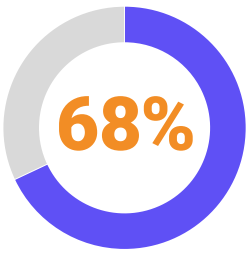

ABANDON PROJECTS

1. Significant time, resources and budgets are being wasted.

While ML practitioners often run, adjust, and re-run experiments as part of model development, **68%** of respondents admitted to abandoning **40-80%** of their experiments due to mismanagement.

Respondents cited issues including lack of infrastructure or resources, reproducibility failures, API integration errors, and debugging failures. These problems are due to breakdowns that occur throughout the machine learning lifecycle outside of the typical iterative process of experimentation.

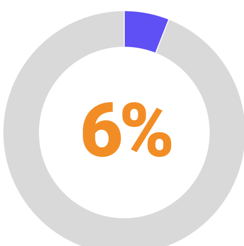

DEPLOY MODELS IN FEWER THAN 30 DAYS

2. There is a lag in model deployment.

Just **6%** of teams surveyed have been able to deploy a model into production in fewer than 30 days. Nearly half **(47%)** of ML teams require four to six months to deploy a single ML project, while another **43%** take up to three months.

The early stages of ML model development are iterative and experimental, driven by data science frameworks and processes. When data science teams hand off trained models to engineering teams to rebuild and/or deploy in production environments, issues can arise that affect model viability and performance. All of these factors can delay model development and deployment, hampering teams' ability to deliver value to their respective lines of business.

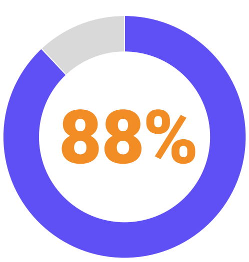

HAVE ANNUAL BUDGET OF LESS THAN $75K

3. Budgets for tools that could address issues are woefully inadequate.

The majority **(88%)** of respondents have an annual budget of less than $75,000 for machine learning tools and infrastructure. That is less than the average salary for a data scientist, which ranges from $121,000 and $149,000 depending on location, according to Salary.com.

The average annual budget of respondents is dwarfed by the opportunity cost resulting from under-investment in ML tools and MLOps. While there is a growing ecosystem of tools, teams face many choices, including the build-versus-buy consideration: Should they build based on open source tools to add capabilities to their ML development processes or should they buy tools to fill those gaps? Whatever the choices, adequate resources must be allocated to support the work of ML teams.

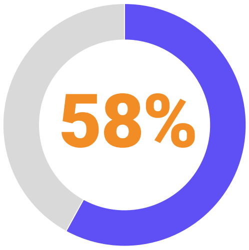

USE MANUAL TOOLS

4. Without funds for automation, ML teams must track experiments manually.

Over half **(58%)** of respondents reported they are using manual tools, such as spreadsheets or Git, to track ML experiments. This can introduce human errors and friction into the process, causing projects to take longer to complete.

Manual tracking also creates challenges for team collaboration, dataset versioning, and model lineage tracking. The downstream effects of these issues can include hindered model auditability and difficulty reproducing proven approaches.

It's important to note that companies are not intentionally withholding budgets or misallocating ML resources.The majority of survey respondents **(63%)** said their organizations would increase machine learning budgets in 2022.

Getting to the Good Part: Less Friction, More ML

Machine learning can deliver outsized business value but the process of developing ML is complex, with many dependencies and potential pitfalls. ML initiatives and techniques are in the early days of development, and as the survey results show, practitioners are managing friction in the form of challenges related to infrastructure, integration, and resources.

But there's good news: companies like Netflix and Uber have pioneered solutions to some of these challenges with excellent results, and some have made their in-house tools available via open source. These lessons, documented in engineering blogs and conference sessions, along with a growing ecosystem of tools, are making it easier for ML practitioners to build better ML models faster.

Recommendations for ML Practitioners Evaluating Tools

There are three characteristics practitioners should consider when evaluating tools to reduce friction and accelerate the ML development process:

1. **Integration and customizability:** Look for tools that complement effective processes already in place. The best tools are highly customizable to practitioners' workflows and integrate seamlessly with proven frameworks.

2. **Scalability:** Any company investing in ML teams and infrastructure should ensure the tools they work with can handle the intense demands of massive scale. To better understand capacity, practitioners can ask tool providers about scalability by the numbers: models managed, concurrent training jobs, and datasets stored.
3. **Complete ML lifecycle:** ML development is iterative and cyclical. For example, when models fail in production, it's helpful to revisit their performance baselines in training and retrain them as appropriate. Practitioners will want tools that make it easy to create and visualize those comparisons across the complete ML lifecycle – from data preparation to monitoring models in production.

ML promises a virtually limitless opportunity to deliver value. ML practitioners who embrace new tools and techniques will accelerate model development and get to the good part faster: realizing value for their customers, teams, and companies.

About Comet

Comet provides a self-hosted and cloud-based machine learning platform for enterprise, academic, and individual teams. Data science and ML practitioners use it to track, compare, explain, and optimize their models across the complete ML lifecycle – from training runs and model management to monitoring models in production.

Comet's highly scalable ML development platform is trusted by over 150 enterprise customers including Cepsa, Etsy, Uber and Zappos. Comet is free for individuals and academic teams. Startup, team, and enterprise licensing is available.

To learn more, visit comet.ml or contact us.