

AI That Helps Data Scientists Guide Organizations from Turbulence to Recovery

As businesses navigate the difficult path from disruption to recovery, the role of the data scientist is more important than ever.

With most data either out of date and new data changing erratically, organizations are counting on their data scientists to guide them with careful scenario planning, what-if analysis and rapid AI development, with models able to adjust quickly to changing business conditions.

Aible is a collaborative AI suite that enables data scientists to help their organization make critical decisions at a time when businesses are being challenged as never before. Aible delivers a rapid path to production and gets models into enterprise applications quickly so end users can provide feedback, which enables businesses to adapt quickly to changes in the market. At the same time, Aible locks down best practices from data to modeling decisions and delivers business impact without ever compromising security, so data scientists are assured of complete data control at all times.

Aible was developed by listening to data scientists and creating an AI suite that responds to four of their biggest challenges:

1. Data Scientist challenge:
“Models keep getting stale and I have to keep retraining the same AI instead of working on new projects.”

Unique among AI platforms, Aible automatically deploys an efficient frontier of dynamically-balanced models rather than a single – and often fragile – model. Aible’s portfolio approach means that as market change, other models that have been pre-calculated to reflect those new business realities can be used instead of always having to retrain models. Aible gives data scientists a rapid path to production, a critical advantage in turbulent times. Integration with end applications allows real-world monitoring to show business impact as well as “model drift.” Aible monitoring also collects the necessary data for retaining and automatically informs the user when sufficient new data has been collected so retraining is possible.

2. Data Scientist challenge:
“I can’t prove the value of my models because they are not getting into production and adopted.”

Aible builds trust with business users by making them an important part of the AI process. Aible empowers business users and analysts to actively contribute to many aspects of the machine learning process. Scenario analysis shows users how the model would affect their business. Rapid prototyping gets models working in applications like Salesforce so end users can provide feedback. Aible ensures that users understand exactly how models affect their KPIs without them having to understand model metrics or constantly consult the data science team for minor adjustments.


3. Data Scientist challenge:
“I have to do similar projects over and over again.”

Data Scientists can empower others to train models by building Blueprints that encode best practices for creating the right AI for common use cases, from data to modeling decisions. So, if they have created a model for one geography or one product, they can quickly use the blueprint to create custom models for all other geographies or products. Other users can also use the Blueprints as starting points for their own projects that are similar to the Blueprint created by the data scientist.

4. Data Scientist challenge:
“How do I know my data is secure?”

Aible ensures data security by never accessing user data. Aible runs directly in the organization’s secure AWS or Azure cloud account, or on-prem, ensuring complete data security. All data stays in the secure cloud account, where the Aible models are trained and deployed. All data transformation, training, and testing code are logged in the account. Aible never takes possession of company data, and can’t see any of the models. It also maintains an audit trail of all decisions and changes made by every user.

Data scientists have a familiar toolbox full of statistical metrics that they use to train AI models, everything from Recall, Precision and F1 to AUC, Log Loss, and Mean Absolute Error. A junior data scientist fresh out of school who suddenly gets thrown into the business world will invariably start by selecting what they think is the best metric on which to train a model, given the business context. But experienced data scientists know that approach often doesn't deliver the impact the business is seeking.



Standard Optimal Curve

Rather than pre-selecting a single metric on which to train an AI model, adjusting it for business context, and hoping it aligns with the needs of the business, it's far more effective to deploy a portfolio of many models that are optimized for a wide range of business realities. Unique among AI providers, Aible deploys what's known as an efficient frontier of models rather than a single model, so that as conditions change, businesses always use the right model at the right time. The efficient frontier enables businesses to spread their bets with AI models – a significant advantage in uncertain times, when making a single bet is riskier than ever.

Many model metrics work better in theory than in the real world of business. Whenever you train a model on a single metric, it's going to be optimized specifically to that statistical measurement, not necessarily to what the business wants or needs. Statistical metrics have a way of satisfying mathematics, but not necessarily CEOs.

Data scientists have long recognized that a single statistical metric rarely aligns perfectly with what the business needs so they perform threshold adjustment. When you train a model on a single metric that may or may not fit the business in the first place, it's difficult to then adjust it in a satisfactory way. Sometimes, the orthodox method of model training and adjustment ends up creating the worst-performing model in terms of business impact. A poorly-performing AI model may be worse than useless – it can actually destroy business value and result in the loss of millions of dollars. For example if the success rate for a sales process is relatively low, then models trained on precision or accuracy may be too conservative, tell salespeople to pursue too few deals, and thus destroy economic value. When that happens, it's no consolation to the business that the model's accuracy score is really high.

That's why Aible's portfolio approach is better for business. Aible deploys an efficient frontier of many models trained to different business assumptions, rather than one model tuned to a single metric and adjusted after the fact. When business conditions change, Aible lets you move along the efficient frontier and select another pre-calculated model that's in tune with the current business realities, rather than having to do a heavy-lift model retraining. The right model will change over time as business conditions and capacity constraints change. The best model isn't the one that scores highest on a statistical metric, it's the one that delivers the greatest business impact – more revenue, lower costs, fewer customers lost to churn, optimized sales and marketing campaigns, more efficient supply chains, less waste and fraud. At every step, Aible monitors predictions and recommends model updates. That's a critical advantage in uncertain economic times, when business assumptions are changing rapidly and conventional AI models are decaying quickly.

Of course, there are instances in which a model trained to a single statistical metric may in fact be exactly right – just as a broken clock tells the correct time twice a day. If your business happens to line up perfectly with a recall metric, then recall is definitely the way to go.

But in reality, businesses rarely align perfectly with a point on the broken clock. Particularly in turbulent economic times, a model trained to a single statistical metric is almost always wrong, like a broken clock. With the efficient frontier, Aible can tell you the right time anytime of the day.

Statisticians like to categorize errors into Type 1 (False Positives) and Type 2 (False Negatives). But experienced data scientists know that the most important errors in AI are “Type 3” errors. The term was coined in 1948 by statistician Fred Mosteller, the founding chairman of Harvard’s statistics department. Mosteller defined a Type 3 error as “correctly rejecting the null hypothesis for the wrong reason.” In colloquial terms, a Type 3 error means giving the right answer to the wrong question.

“Most AI companies are built to get data scientists predictions. Aible is built to get them promoted.”

With the commercial adoption of AI, much time, money, and effort has been spent building mechanisms to fit bigger and bigger models using larger and larger data centers. For a certain class of problems this approach is the right answer. For many others, it is a Type 3 error. The problems that demand massive data infrastructure tend to be those where the goal is either highly personalized service, or the need to classify items from a large catalog. For many more traditional data problems there are diminishing returns to huge models. Despite the tremendous resources devoted to AI infrastructure, the main problem with AI and data science today isn’t that the models are not good enough. It’s that the models are not getting used.

A 2019 MIT-Sloan/Boston Consulting Group survey found that 65% of companies reported that they are seeing little or no value from the AI investments they’ve made in recent years. Nearly 40% of organizations that have made “significant investments” in AI say they have not seen any business gains. The reality is that many companies are spending enormous amounts of time and money on AI and seeing little or no payoff. AI models that take many months to develop often wind up never seeing action in production. Without business adoption, they have produced no value.

Aible takes a different approach. Aible provides a forum where data scientists and business users collaborate as a team. Business users are an integral part of the AI process; they have a say in setting assumptions, providing the

business expectations, and giving feedback about the models as circumstances change. This alleviates the problems data scientists commonly have in trying to explain the results of the modeling and have the business people understand it.

Frontline data scientists have to deal with modeling complexities that business users never see and wouldn’t understand if they did – choosing which variables to include and which to leave out, testing the model to make sure it’s valid, making sure the model doesn’t violate any regulatory restrictions. A common reaction to that process by business people is: Why is it taking so long?

Aible removes that friction by automating much of the rote part of the model fitting process, by simplifying the administrative work around “owning” a model, and by translating model metrics into language business people understand – dollars and cents. What drives business adoption of AI is the promise of more revenue and lower costs, not a high F1 score. Aible makes it easy to go straight from the modeling step to deploying the model and having it used in production. It also automates the monitoring process to signal when models need to be adjusted.

Aible optimizes models for business impact, with the incredibly useful byproduct that the impact gets measured in the first place. This is obviously useful for business leaders looking to justify investments in their data science teams, and it is a veritable gold mine for data scientists.

Ironically, the people who are among the most quantitative in their organization often struggle to put a dollar value on their work. This can put data scientists at a disadvantage during performance reviews. When raises are determined by contributions to the bottom line, how should that improved F1 score be judged against a sales or marketing initiative that brought in an easily understood \$X million dollars? Aible automatically quantifies a data scientist’s contributions in the financial terms best understood by their business stakeholders.

Conventional AI struggles with adoption because of a fundamental Type 3 error. Too much time has been spent precisely answering questions about how to make better models, while ignoring the high probability that the models won’t be used in the first place. Aible is built from the ground up to aim at the right target, to get models adopted, and to deliver value to the business.