Addressing the Scale-up Challenge through Direct Impact

What is Georgian's reading of the scale-up challenge?
Typically, software companies at the scale-up stage have established product-market fit, but must overcome a variety of barriers to growth. For the scale-up ecosystem as a whole, these barriers can include access to capital, talent, and markets. For SaaS companies with data-driven products, these barriers can include meeting data privacy and security requirements, optimizing machine learning processes, establishing data feedback loops, and earning trust with users to drive product adoption.

What is Georgian's response to address the challenge?
Georgian takes a thesis-driven approach to investing, specifically identifying companies that can gain momentum and overcome the scale-up challenge by taking advantage of universal technology trends including Applied AI, Security First, and Conversational Business. To this end, Georgian has made considerable investments in data, people, and technology to enable Georgian to directly assist companies in these areas post-investment.

How does Georgian help in scaling up companies?
Georgian's value-add approach is executed through the Georgian Impact team, which includes experienced technology practitioners with training in statistics, machine learning, software engineering, linguistics, and natural language processing. The Impact team helps the technical staff of companies strengthen their competitive advantage by identifying and prioritizing opportunities to incorporate thesis work into company strategy and accelerating product development through applied research, capacity building, rapid prototyping, and consultation.

Currently, the research roadmap focuses on protecting customer data assets, accelerating machine learning processes, and increasing the transparency and explainability of models.

Research Area	What?	Why?
PROTECT	• Differentially private machine learning to enable quantifiable privacy guarantees and support data aggregation, regulatory requirements, and IP protection	• Acquire broad data rights and build a competitive moat through data network effects and IP • Reduce the cost of meeting regulatory requirements, such as GDPR • Minimize data security and privacy risk
ACCELERATE	• Accelerate machine learning model training through the use of simulation, cross-customer data, pre-trained third party models or third party data	• Speed up onboarding cycles and minimize days to revenue and billings of AI software • Improve AI performance and generate higher outcomes
EXPLAIN	• Incorporate explainability and fairness capabilities into the predictions and decisions of AI-driven solutions across all industries according to user requirements and regulation	• Earn trust with users • Enable and increase AI software adoption • Meet regulatory requirements and minimize business risk

Georgian Partners Applied Research Areas

What are the initial results?
The Georgian Impact team has completed a number of applied research projects that have resulted in new portfolio product launches and improvements to production systems. The attached case study explains how differential privacy was applied to gain lift in predictive models on aggregated data while maintaining measurable privacy guarantees for customers.

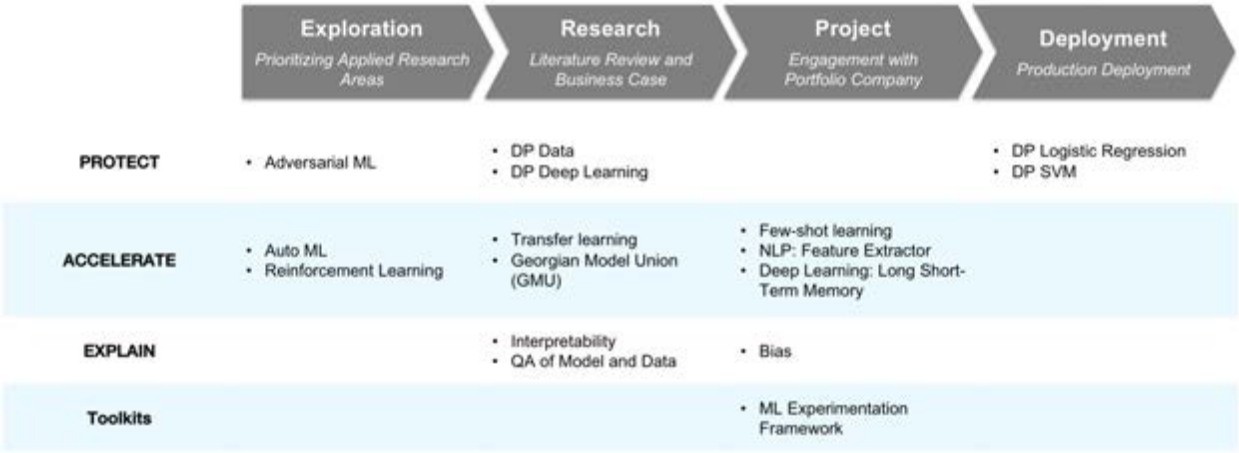

Georgian Partners Applied Research Roadmap

What impact could this model have on the greater scale-up ecosystem?
Georgian is exploring a number of extensions of this model to benefit community members beyond the portfolio. These include, but are not limited to:

- Establishing data unions and data exchanges
- Sponsoring academic research partnerships
- Publishing scientific papers
- Open-sourcing code libraries, models, and data sets

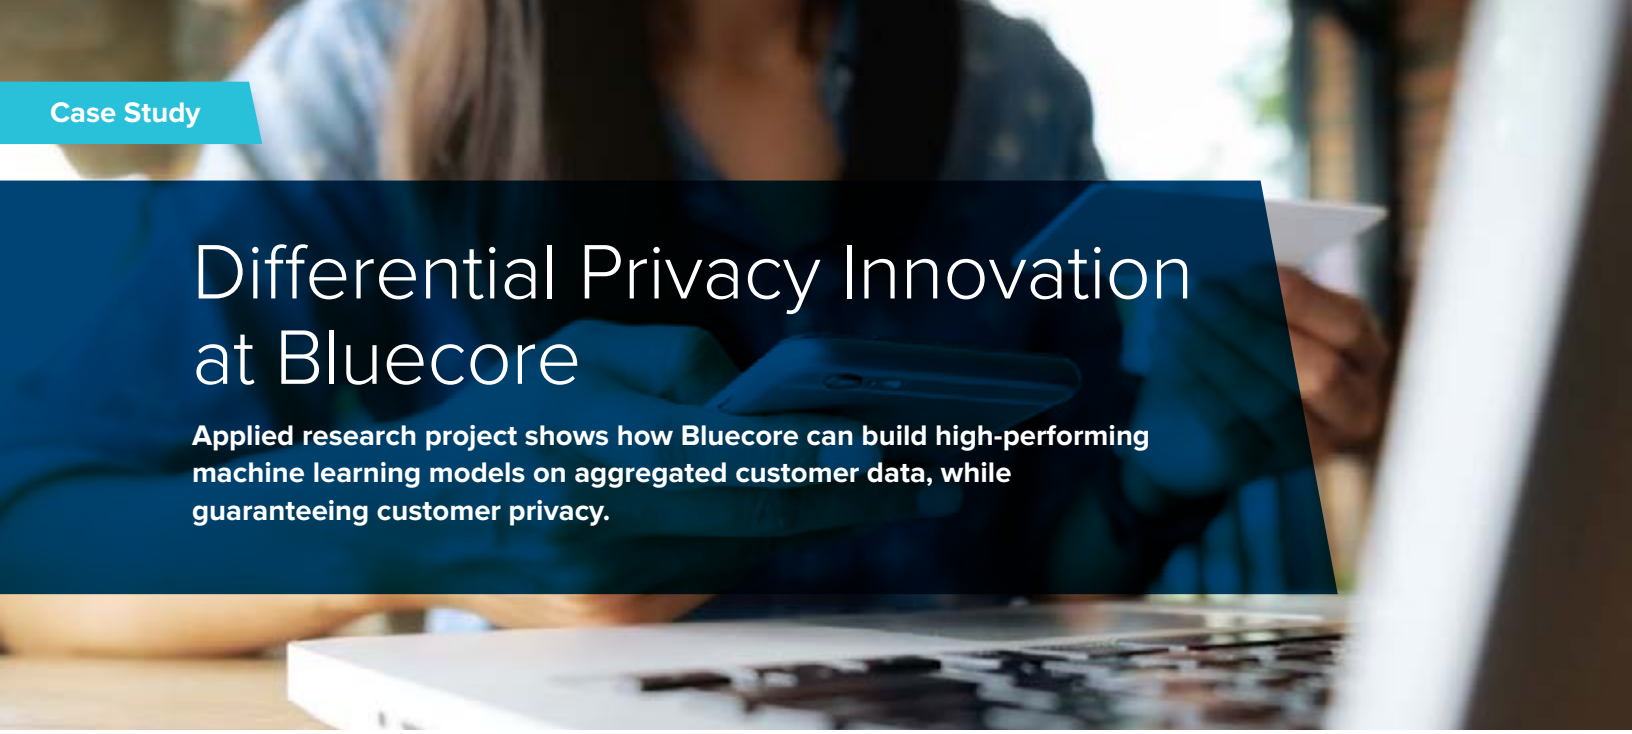

Case Study

Differential Privacy Innovation at Bluecore

Applied research project shows how Bluecore can build high-performing machine learning models on aggregated customer data, while guaranteeing customer privacy.

As the leading retail marketing platform specializing in email, New York City-based Bluecore uses data to help companies identify their best customers and keep them for life. It does this by combining customer behavior and product data, which allows marketers to create personalized campaigns across all their marketing channels. Additionally, Bluecore's retail-specific predictive models can help retailers find which customers would be most likely to buy at a given time, which customers are likely to unsubscribe from their email list, and which products will be most likely to drive a second purchase.

It's Bluecore's add-on services — Propensity to Convert and Product Affinity — that are among its most popular and sought-after offerings. However, as with all B2B SaaS companies that provide data-driven insights, Bluecore has long faced what is known as the 'cold start' problem. Specifically, these offerings work best once sufficient data has been collected from a new customer to build an accurate predictive model for that customer. Collecting that data can take anywhere from several weeks to six months.

This challenge meant that Bluecore could only offer its analytical services as an upgrade to existing customers, rather than as part of their core offering at the outset of each new customer relationship. That translated into a longer time to value for new customers and delayed revenue opportunities for Bluecore.

Bluecore knew that the only way to get around this cold start issue was to build a predictive model that used all of its customers' data collectively. That way, new customers could instantly benefit from the model, even if they didn't yet have enough data of their own. Unfortunately, given the high level of sensitivity around data privacy, this never seemed like a viable option. That's because most customers are concerned that using their data to build models for others could result in leaked trade secrets or transactional information to their competitors.

Differential Privacy Offers a Solution

After leading Bluecore's Series B round of funding, Georgian Partners believed they had a way to solve the company's cold-start problem: differential privacy. In late 2016, the Georgian Impact team pitched Bluecore on how differential privacy might allow the company to deliver both privacy guarantees and better insights.

Differential privacy provides a mathematical definition for the privacy loss that results to individuals when their private information is used to create a data product. This makes it possible to mathematically show how private a machine learning model or query function is. While Microsoft Research pioneered the original work on differential privacy, it's been Google and Apple that have recently brought differentially private products to market.

Bluecore Co-Founder and CTO, Mahmoud Arram, was immediately intrigued with the idea. While familiar with differential privacy conceptually, he knew that it wasn't a mainstream practice and certainly not something that many other growth-stage software companies were doing. "It was definitely a moonshot project for us," recalls Arram. "It was a risky move, but also the chance to do something very innovative. And, if we got it right, I knew that the rewards would be great because we'd be positioned to do something that our customers couldn't do on their own."

After brainstorming ideas, Bluecore and the Impact team got to work trying to figure out how they could use differential privacy to solve the company's challenges. They worked intensely over the weeks that followed, pouring over research and sample customer data that Bluecore had been given permission to use.

The combined Bluecore and Georgian Partners project team tried several different approaches and combinations of machine learning models and differentially private transformations. At first, the efforts didn't yield usable results. But, the team persevered and eventually discovered a combination of machine learning models, optimization algorithms and differentially private transformations that produced both high accuracy results and a high degree of privacy (ε= 0.01 in differential privacy terminology).

"That work paid off," says Madalin Mihailescu, Georgian Partners' CTO. "Using the data sets that we had, the team was able to create a differentially private model that allowed Bluecore to dramatically increase the accuracy of its predictions around conversion."

The project has shown that both new and existing Bluecore customers can benefit from the use of differential privacy to enable data aggregation. New customers will be able to immediately start predicting which website visitors are more likely to buy from them, even if they don't have their own data to support it. Meanwhile existing customers are predicted to benefit from a 10 percent increase in sales, representing a significant uplift.

After their initial success with applying differential privacy at Bluecore, the Impact team has been working with the company to productize the code and has made the first version of its differential privacy product available to its wider portfolio. "The Impact team made a tremendous contribution to our business," says Arram. "Thanks to their input, we now see differential privacy as one of our main points of differentiation."

Bluecore plans to roll out differential privacy in its Propensity to Convert and Product Affinity offerings in a future release. For Mihailescu and team, that's the ultimate validation. "It's great to see our applied research and product development efforts help our portfolio create real value for their customers," he says.

> "The Impact team made a tremendous contribution to our business. Thanks to their input, we now see differential privacy as one of our main points of differentiation."

Mahmoud Arram,
Bluecore Co-Founder and CTO