Riskfuel

1,000,000x faster models

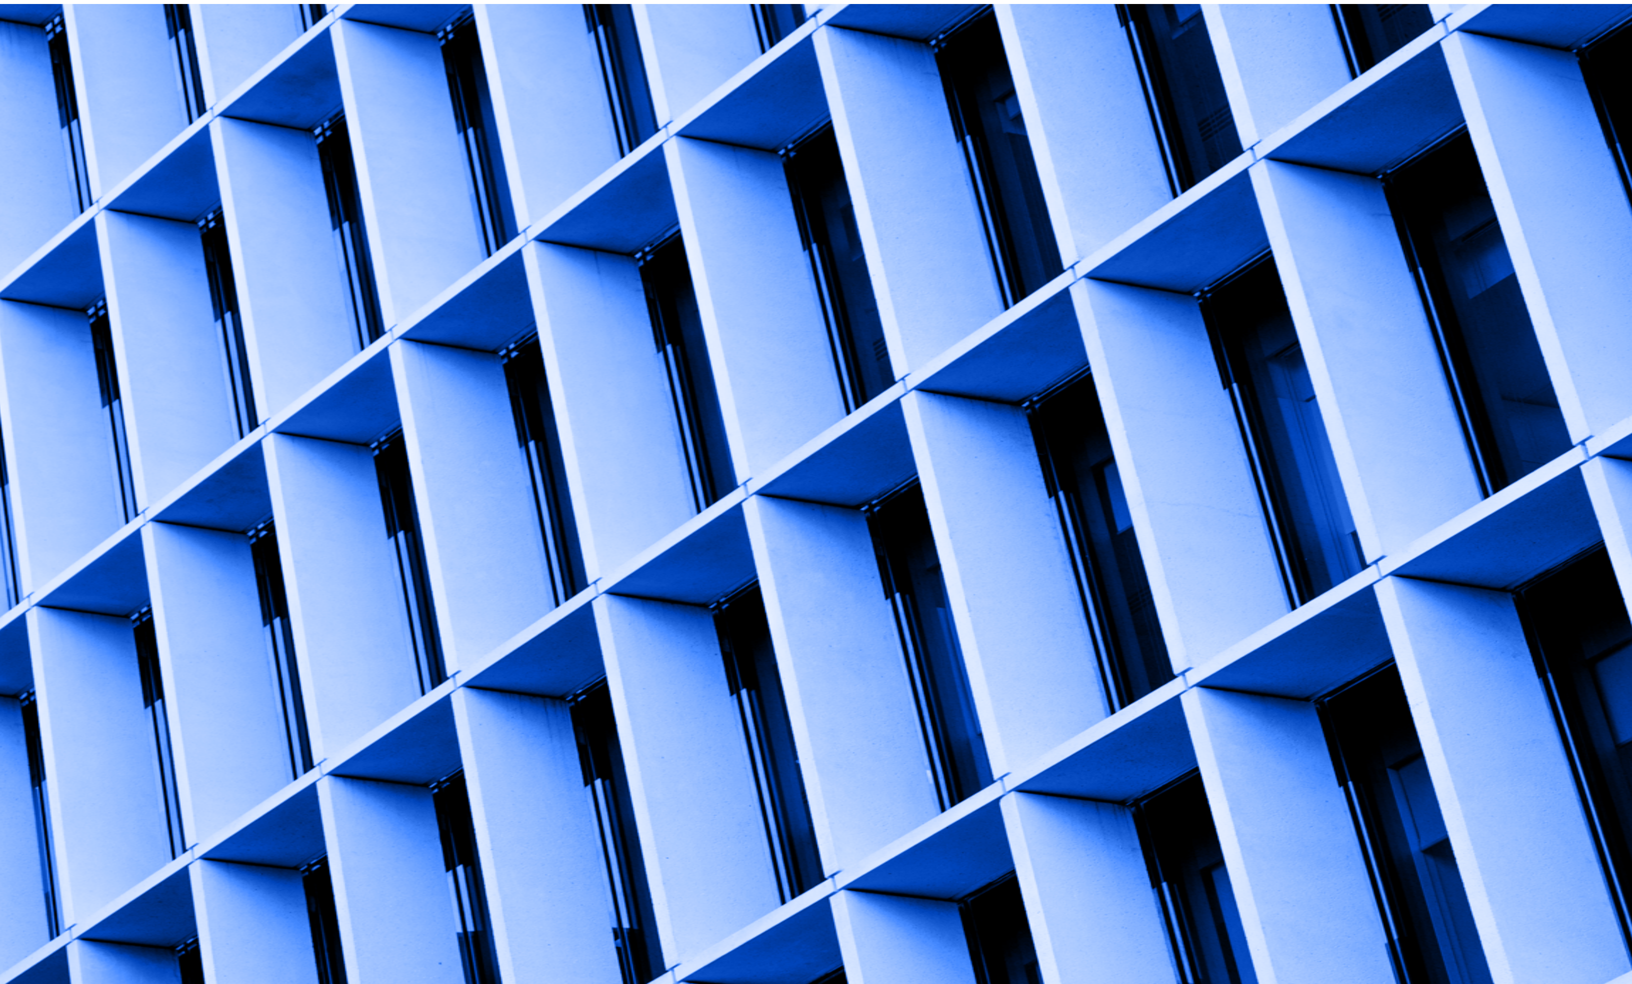

1,000,000x faster models: how it works

Traders of structured financial products are held hostage to overnight batch processing. Their risk sensitivities are available to them only once a day, making efficient real-time risk management of their positions immensely difficult. Riskfuel AI-based technology cuts the computation costs to virtually zero allowing for on-demand recalculation of portfolio values and a complete up-to-the-second view on risk.

Problem

Exchange-traded products have seen a tremendous wave of technological innovation in the past 20 years, leading to much more efficient markets and dramatically lower transaction costs. Such innovation included switching to decimal quoting, electronic trading, and real-time data streaming. With instantaneous price discovery allowed by exchanges, cash traders are aware of their portfolio valuations and risk sensitivities in real time.

The situation is very different for structured financial instruments traded over-the-counter (OTC). The true market price for a derivatives contract (a swap, an option, or an exotic derivative) is only known when the contract is signed and when it is sold or finally matures. During the life of the contract in a trader's book, mathematical models must be deployed to value the transactions and provide the trader with basic tools to manage the risk in their portfolios. Such mathematical models are typically built by teams of financial engineers in the front office and risk management departments. The modelers have made a lot of progress in optimizing the traditional computational techniques, such as numerical integration and Monte Carlo simulation.

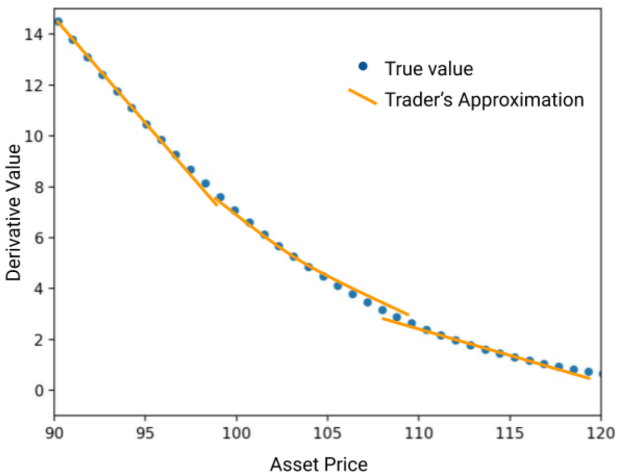

Typical quality of numerical approximation in a front-office tool used during trading day. The approximating function is calculated at a few points and extrapolated using linear or quadratic functions. The problem becomes more pronounced in many dimensions due to the curse of dimensionality.

However, **the goal of real-time evaluation and risk management remains elusive.** Each step forward towards computational efficiency (Moore's Law, Cloud, parallel processing, better algorithms, algorithmic differentiation) is met with demands for even more computations (more valuation adjustments to account for credit risk or funding, more accuracy, more risk sensitivities, new structured products). Throughout the trading day the traders must rely on their simplistic desktop tools, which provide only a rough approximation to production models. Even the basic question of "what's my P&L today" only gets answered the next morning.

Deep Neural Networks

Behind the recent success of artificial intelligence is a machine learning technique called Deep Learning. It is based on deep neural networks (DNNs), which are known as universal function approximators. Due to the flexibility of their architecture and an enormous number of parameters (tens of millions to billions in industry-strength models) these neural nets can approximate very complicated multidimensional functions, like a mapping from the matrix of pixels in an image to the text that should describe it. DNNs are particularly suitable for problems where a lot of input and output data is available, and the data remains relevant over time. This is the case in the fields of computer vision, machine translation, and games, where Deep Learning has been tremendously successful.

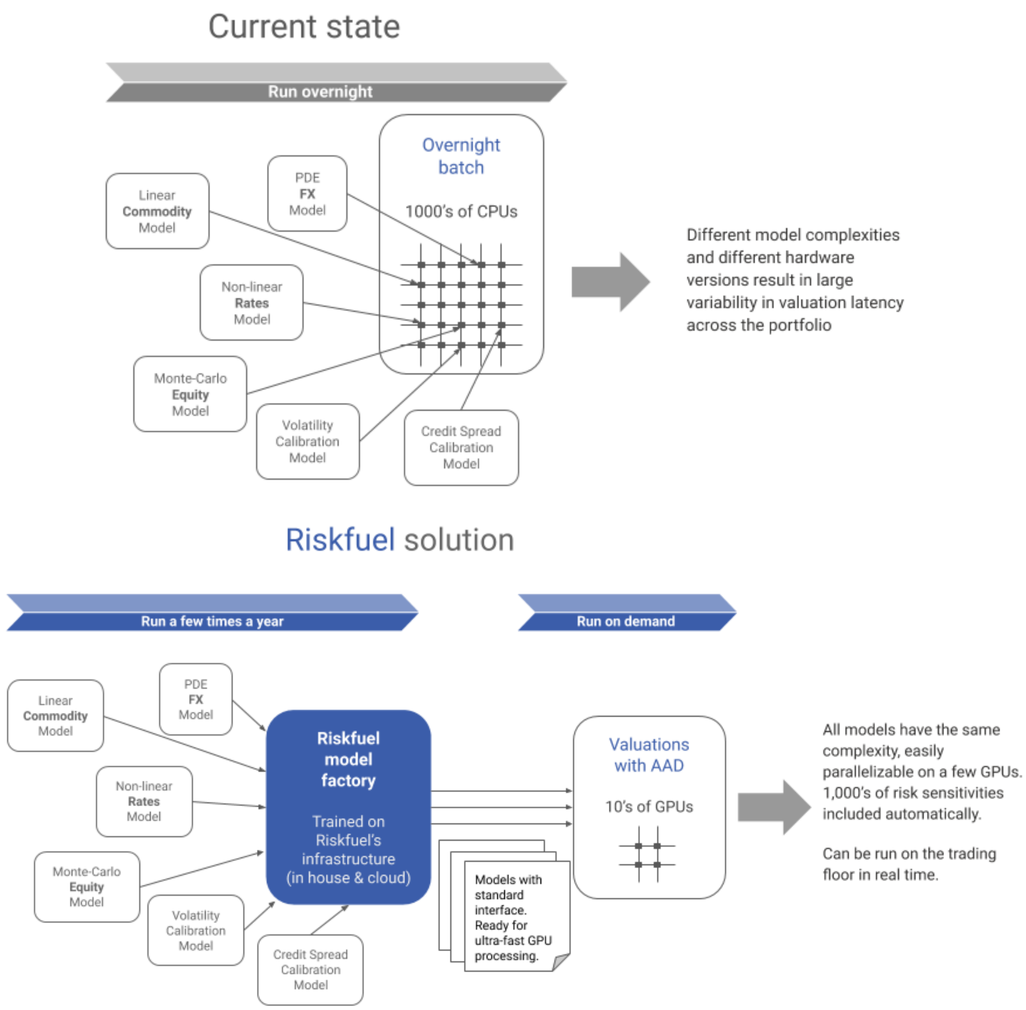

Breaking free from the overnight batch. Top: Current derivative portfolio valuation process. Bottom: The process amended by Riskfuel technology. Model training and update is maintained by Riskfuel. We support both on-premise and Cloud deployment.

Riskfuel uses existing derivatives valuation models to generate training data for its DNN approximators. The valuation models can be industry-standard or customer-specific. This generated data is available as needed and can be obtained in unlimited amounts if required for precision and fine tuning. Our proprietary methodology explores the model parameter space most efficiently within a large range of market conditions. The relatively high compute time cost for data generation and training is incurred infrequently (a few times a year). The payoff during the trading day, when the speed and accuracy matter most, is tremendous. Riskfuel's DNN-based pricers produce on-demand valuations and risk sensitivities in milliseconds, up to a million-fold improvement compared to traditional models.

While DNNs provide an extremely efficient way to represent complex pricing functions calculated by sophisticated exotics valuation models, they also help speed up calculations of simpler instruments such as forwards, swaps, and vanilla options. That is because encoding even simple models in neural network forms allows one to take advantage of massive parallelization capabilities afforded by modern deep learning technology and specialized GPU and TPU hardware. Moreover, since gradient calculations are the heart of the neural network training process, all risk sensitives are calculated automatically at no extra computational cost for development or runtime.

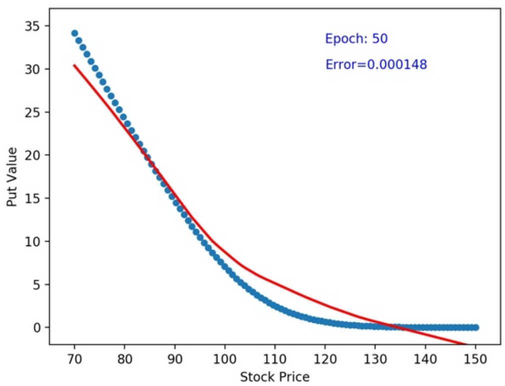

Early stage of training

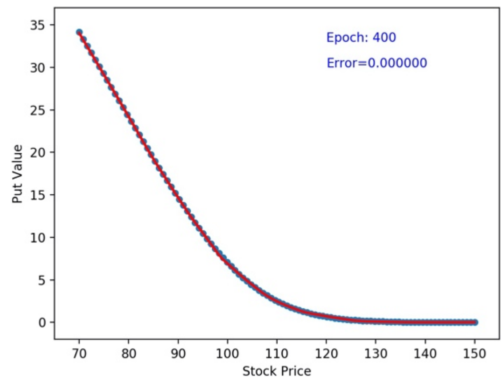

After training

A simple example of DNN training on a put option value function. The blue dots represent training data generated by traditional models. The red line in a DNN approximation. After training, DNNs become smooth, very accurate and very fast approximators.

Model Risk

Accuracy

Riskfuel models are very accurate. Our training data is generated from the client's production models. We work with the models developed by the client's financial engineering teams internally as well as industry-standard models provided to clients by vendors. Riskfuel's models typically end up being more accurate than the target benchmark model for the following reasons:

- Since our training process is done very infrequently, we can afford to generate training data at very high levels of accuracy that would be overly expensive for regular production use.
- Properly regularized DNNs do not "overfit" the noisy data, they provide smooth approximations instead.

Model validation

Since during the training process Riskfuel models are essentially benchmarked to the client's production models, the resulting DNN approximators can be easily validated by running our models against the models in production. Due to the fast speed of Riskfuel pricers, they can be quickly stress-tested under a large set of market conditions within the training domain. We work with your Model Validation groups to satisfy any additional requirements they may have.

Retraining

Riskfuel models are built to perform within a wide range of market conditions. We commit to periodic retraining of our models to accommodate any major market movements.

Frequently Asked Questions

I want to know more contact us at: sales@riskfuel.com

More technical readers may want to review our technical paper *Deeply Learning Derivatives* by R. Ferguson and A. Green (October 2018), SSRN, also available at https://arxiv.org/abs/1809.02233

You claim a 1,000,000x speed-up. Really?

Yes. While the level of improvement depends on the accuracy settings in your production models, a million-fold speed-up is indeed achievable. This is due to a combination of factors:

- Fast DNN approximators instead of slow traditional models
- GPU instead of CPU hardware
- Fewer number of processors, avoiding network bottlenecks

Still skeptical? Check out how a group of scientists achieved a similar feat solving a 300-year old partial derivatives equation: *A neural net solves the three-body problem 100 million times faster*, MIT Technology Review, October 2019

Your paper talks about a call on basket of equities, but I want an FX digital one-touch knock-out!

We've been busy since our first paper! We have implemented a range of models in various asset classes: interest rates, FX, equities. To put it simply, if you can solve it using a traditional financial engineering model, we can encode it into an ultra-fast DNN.

What about sensitivities?

Modern deep learning technology provides automatic ways for differentiating models. You get sensitivities to all model inputs for free, whether you have implemented Automatic Algorithmic Differentiation (AAD) or not.

Can this solve all my problems, like XVA or FRTB?

Significant speed-up will benefit many processes that depend on derivatives pricing models, including XVA, FRTB, market and credit risk limits, stress testing, capital and liquidity management.

We have been building our production processes for 20 years. How do we integrate your models?

We realize that onboarding any new vendor is a journey. We make it as easy as possible for you. The interface to our models is client-specific. The Riskfuel pricers use the same inputs and produce the same outputs as your production models.

Sounds cool. We're already looking into it. What do you bring to the table?

Unmatched expertise spanning both capital markets and artificial intelligence. Our Founder and CEO had a 15-year career as a senior quant. He pioneered the field of Deep Learning for derivatives valuations in his paper cited above. Our CPO had a 10-year career as a quant and risk manager, and he spent 4 years in Data Science and AI.

We are in the middle of migrating to Cloud. When would be a good time to call you?

Call us now! Riskfuel technology integrates seamlessly with Cloud as well as on premise. We use modern technologies that ensure easy deployment in either scenario.

Can you give us a demo or a proof of concept?

With pleasure! Please contact us at sales@riskfuel.com