Cutter Consortium
Access to the Experts

Cutter Business Technology JOURNAL

Vol. 30, No. 4, April 2017

Management, Innovation, Transformation

"While much of the emphasis in fintech R&D remains on emerging technologies, the industry is also looking at how to efficiently integrate various technologies, platforms, and systems together."

— Philip O'Reilly, Guest Editor

The Frontier of Fintech Innovation

Cutter Business Technology JOURNAL

As business models for creating value continue to shift, new business strategies are constantly emerging and digital innovation has become an ongoing imperative. The monthly *Cutter Business Technology Journal* delivers a comprehensive treatment of these strategies to help your organization address and capitalize on the opportunities of this digital age.

Founded in 1987 as *American Programmer* by Ed Yourdon, and until recently known as *Cutter IT Journal*, *Cutter Business Technology Journal* is unlike academic journals: Each monthly issue, led by an expert Guest Editor, includes five to seven substantial articles, case studies, research findings, and/or experience-based opinion pieces that provide innovative ideas and solutions to the challenges business technology professionals face right now — and prepares them for those they might face tomorrow. *Cutter Business Technology Journal* doesn't water down or delay its content with lengthy peer reviews. Written by internationally known thought leaders, academics, and practitioners — you can be certain you're getting the uncensored perspectives of global experts.

You'll benefit from strategic insight on how the latest movements in digital innovation and transformation, IoT, big data analytics and cloud, to name a few, are changing the business landscape for both new and established organizations and how cutting-edge approaches in technology leadership, enterprise agility, software engineering, and business architecture can help your organization optimize its performance and transition to these new business models.

As a subscriber, you'll also receive the *Cutter Business Technology Advisor* – a weekly bulletin featuring industry updates delivered straight to your inbox. Armed with expert insight, data, and advice, you'll be able to leverage the latest business management thinking to achieve your organization's goals.

No other journal brings together so many thought leaders or lets them speak so bluntly — bringing you frank, honest accounts of what works, what doesn't, and why. Subscribers have even referred to the *Journal* as a consultancy in print and likened each month's issue to the impassioned discussions they participate in at the end of a day at a conference!

Get the best in thought leadership and keep pace with the technologies and business models that will give you a competitive edge – subscribe to *Cutter Business Technology Journal* today!

Cutter Business Technology Journal®

Founding Editor: Ed Yourdon
Publisher: Karen Fine Coburn
Group Publisher: Christine Generali
Managing Editor: Karen Pasley
Production Editor: Linda Dias
Client Services: service@cutter.com

Cutter Business Technology Journal® is published 12 times a year by Cutter Information LLC, 37 Broadway, Suite 1, Arlington, MA 02474-5552, USA (Tel: +1 781 648 8700; Fax: +1 781 648 8707; Email: cbtjeditorial@cutter.com; Website: www.cutter.com; Twitter: @cuttertweets; Facebook: Cutter Consortium). ISSN: 2475-3718 (print); 2475-3742 (online).

Subscription rates are US $485 a year in North America, US $585 elsewhere, payable to Cutter Information LLC. Reprints, bulk purchases, past issues, and multiple subscription and site license rates are available on request.

Opening Statement

by Philip O'Reilly, Guest Editor

It's a pleasure for me to introduce the first of two special issues of *Cutter Business Technology Journal (CBTJ)* showcasing the thought leadership and cutting-edge research and development (R&D) being done in State Street Corporation's Advanced Technology Research Centres in Europe, the Middle East, and Africa (EMEA) and Asia Pacific (APAC), in partnership with University College Cork (UCC) and Zhejiang University (ZJU), respectively. The articles in this issue represent a small sample of the output from the R&D undertaken in these centers, which combine academic excellence with real industry impact.

I want to begin by providing some context and some insights as to how academia and industry can collaborate in a sustainable manner. The collaboration between State Street, UCC, and ZJU offers a successful model of how institutions of higher education and business organizations can successfully conduct R&D that benefits both academia and business.

Having reviewed various academic/industry collaboration models, I have observed that many extant models struggle to achieve a balance between academic excellence and business relevance. Finding a model that enables all parties to successfully work together for mutual benefit is something that continues to perplex many would-be collaborators. For academics, the primary objective is publishing in highly ranked peer-reviewed journals, while business leaders are primarily concerned with ensuring that the R&D they are funding has real business impact. Therefore, many collaborative R&D initiatives are not sustainable in the medium and long term due to this misalignment between the parties' respective objectives.

The model implemented by State Street in its Advanced Technology Centres in UCC and ZJU is an exemplar in that the products of the R&D efforts include thought leadership pieces in highly ranked peer-reviewed journals and conference proceedings, industry white papers and journal articles (such as those included here), patents, cutting-edge proofs of concept, and software applications that have been implemented by the business. In the case of the papers and articles produced, they are coauthored by academic and industry colleagues. Furthermore, the research efforts have informed and benefited multiple stakeholders in the financial services ecosystem, including international standards bodies, customers, other financial services companies, and academia.

The four articles published in this issue discuss some of the key technologies that will be of significant relevance to the future of financial services and, potentially, other domains. Specifically, they focus on semantic ontologies, next-generation robo-advisors, tools supporting internationalization and localization of legacy systems, and 3D visual analytics.

Financial services organizations must not only stay abreast of emerging fintech technologies, but also effectively integrate them with legacy information systems within an organization's existing IT ecosystem.

While much of the emphasis in fintech R&D remains on emerging technologies, the industry is also looking at how to efficiently integrate various technologies, platforms, and systems together. The financial services industry is a sector where systems integration is a significant issue, due to ongoing technological developments over multiple decades. Financial services organizations must not only stay abreast of emerging fintech technologies, but also effectively integrate them with legacy information systems within an organization's existing IT ecosystem.

In this dynamic and rapidly evolving ecosystem, efficiently and effectively integrating data from the multiplicity of systems that exist is a significant challenge. Furthermore, CIOs and CTOs must continuously wrestle with the dilemma of when to invest — and in what technologies. Therefore, they face ongoing issues pertaining to standards and systems integration. This is where semantic technologies can play a key role.

Through applying semantic technologies and adding a semantics layer to an organization's IT architecture, the organization's data can remain in its existing system (legacy information system, data warehouse, database, Excel file, etc.) and its existing format (structured, unstructured). By using triple store database technology, data can be transformed into RDF graphs, an industry standard model for data interchange. A critical advantage of this approach is that an organization does not need to copy the entirety of the data universe into a triple store database; only data necessary for the specific use case is transformed.

This approach creates a map of an organization's data, encompassing its characteristics — specifically its data types, properties, and interrelationships — together with an explicit specification of the intended meaning of the vocabulary. Therefore, the semantic data model, through associating meaning with each piece of data, enables faster, better data analysis. The critical advantage of semantic ontologies is not just that they enable data to remain in its host system, but that they also support data standards, facilitating much greater levels of interorganizational data interchange. Furthermore, semantic technologies can play a vital part in dynamically identifying anomalies in an organization's data set through the utilization of inference engines. Among other things, this enables much better data quality for regulatory reporting.

The semantic ontologies approach has huge business value. Critically, when the next generation of fintech technology comes along, an organization won't have to undertake expensive system migration projects. For example, while semantic ontologies can be applied to existing systems, they will also help to maximize organizational value from distributed ledger (aka blockchain) technology.

UPCOMING TOPICS

Roger Evernden
Leveraging Enterprise Architecture for Digital Disruption

C. Patrikakis
Digital Transformation in the Industrial Sector

Philip O'Reilly
Beyond Fintech: New Frontiers

Due to their potential for disruption, blockchain-based systems are attracting a massive amount of interest and investment. However, blockchain technology in and of itself does not address core issues pertaining to data quality and seamlessly facilitating exchange of data that may exist off chain. Therefore, R&D efforts are now looking beyond blockchain as a standalone technology, focusing on how a semantic layer can sit upon a blockchain-based system. A critical advantage of this approach is that it would facilitate a standard semantics-based data dictionary for multi-organizational data and assist with data quality issues, thereby offering a much stronger business case for the implementation of this technology.

In the context of decision support, robo-advisors have received a great deal of attention for their potential to assist both individual and institutional investors with their investments. Indeed, much of the recent focus has been on how these technologies should be operationalized, be it in a fully automated online model or via a hybrid model, whereby the technology is utilized by investment advisors to better inform their clients. Yet from a usability perspective, most of these systems are limited in that their interfaces are typically two-dimensional and text-based. Looking toward the design of the next generation of robo-advisors, researchers are exploring how big data, analytics, and 3D visualization techniques can be combined and integrated to operationalize greater analytical and predictive value. Not only is it possible to use these systems to track the performance of "star traders" and efficiently communicate and illustrate this using 3D visualization techniques, these systems can also predict the future of investment markets. By combining technologies, such systems will play a critical role in decision support by visually informing investors' investment decisions.

Now let's explore the articles in our current edition of *CBTJ* in more detail.

In This Issue

In the wake of the financial crisis of 2007-2008, calls for further regulation of the financial services sector intensified. But as the authors of our first article — Oliver Browne, Nenad Krdzavac, Philip O'Reilly, Mark Hutchinson, David Saul, Dáire Lawlor, and Daragh McGetrick — observe:

> For financial institutions, regulatory reporting has become something of a jigsaw puzzle — one that must be cobbled together into a coherent picture from several boxes into which the pieces from different puzzles have been put over time, for an audience that will never appreciate the

pain involved in organizing that picture or the time and manpower required to build it.

Assembling this puzzle has often meant using Excel macros to pull data from various silos and aggregate and verify it manually, a process that is prone to errors and omissions. Fortunately, the authors argue, there is a better way. Browne et al. introduce FIBO (Financial Industry Business Ontology), "a standard financial ontology language being developed to enable a common understanding among financial institutions" that would "allow aggregation and comparison of data from all institutions on a like-for-like basis and enable [regulators] to fully grasp counterparty risk exposures." They make a persuasive case that a flexible ontology-based approach to regulatory reporting will benefit financial institutions and regulators alike.

In our next article, Jie Yang, Hanxi Ye, Yadan Wei, and Linqian Bao tell us about a hot topic in the fintech world: robo-advisors. These automated, online portfolio managers are making a splash because they offer users entrée to the investment world with lower fees, lower account minimums, innovative features, and user-friendly interfaces. Managing an investment portfolio with little or no human intervention requires finding ways to determine investor goals, identify and allocate assets in pursuit of those goals, and monitor and rebalance portfolios. Yang et al. explain what it takes for robo-advisors to do what they do, bearing in mind both the human and technical perspectives.

Benjamin Franklin famously said that "in this world nothing can be said to be certain, except death and taxes." Many IT practitioners might be tempted to add "and legacy systems." In our third article, Bo Zhou and Lucy Chen note that "some legacy systems developed between the 1980s and 1990s did not take software internationalization and localization into consideration, yet for various reasons ... are still in service around the world." In an age when software applications are distributed worldwide as a matter of course, lack of multi-language support is unacceptable. Zhou and Chen tell us about a framework and tool supports they and their colleagues used to internationalize and localize a large-scale fund accounting legacy system at State Street Corporation. With the help of their source code ranker and automated code search tool, they were able to complete roughly 80% of the reengineering work automatically. Should you wish to replicate their success, the authors assure us that "this framework could be used for other internationalization and localization projects."

Our final article proves the adage "A picture is worth a thousand words." When trying to analyze large-scale, complex temporal data, columns of figures just won't do. Thus, authors Jerry Cristoforo, Qiao Huang, Zhiyu Peng, and Xiaohu Yang introduce Apsara, an interactive visual analytics (VA) system for multidimensional temporal data. After canvassing the work that's been done in visualization methods, systems, and tools, the authors describe how Apsara presents "large-scale data in a vast universe ... enhanced by color, motion, and sound." They explain how Apsara implements two critical facets of effective VA systems — customizable visualization and interactive design — to yield "insight into patterns, trends, and correlations in the data." These capabilities are called into service in a case study in which they use Apsara to visualize China's mutual fund market, both to compare it to the mutual fund market in the US and to predict the growth of China's QDII (Qualified Domestic Institutional Investor) program. Like the authors that precede them, they promise that "this customizable solution can be adapted to different data models and applied to multiple domains."

The articles in this issue of *CBTJ* offer significant insights on key emerging technologies and how they can be integrated together to create substantial business value. They present operational techniques, frameworks, and models that can be used by organizations and stakeholders in the financial services sector and beyond. Finally, they illustrate the value of the R&D model that State Street has implemented in partnership with University College Cork and Zhejiang University, which has resulted in high-impact R&D on the cutting edge of technological innovation in financial services. We hope you will enjoy this journey to the fintech frontier.

Philip O'Reilly is a Senior Consultant with Cutter Consortium's Business Technology & Digital Transformation Strategies and Data Analytics & Digital Technologies practices. He is also Director of the State Street Advanced Technology Centre at University College Cork (UCC), the Academic Co-Director of Boole Business Labs, and a Senior Lecturer in Business Information Systems at UCC. Dr. O'Reilly received the Stafford Beer Medal in recognition of the most outstanding contribution to the philosophy, theory, and practice of information systems from the Operational Research (OR) Society at an awards ceremony that took place at the Royal Society in London. He has been invited to deliver keynotes and workshops by numerous multinational companies and at leading practitioner events, including the Banking & Payments Federation of Ireland National Conference. *Dr. O'Reilly has been a leading member of research teams that have been successful in securing in excess of €2 million in research funding in recent years. His research concentrates on digital business, with a specific focus on financial services. Dr. O'Reilly's work has been published in leading journals in the information systems field, including* European Journal of Information Systems, Journal of Strategic Information Systems, *and* Information Technology and People. *He can be reached at poreilly@cutter.com.*

Enabling More Efficient and Flexible Reporting Through Semantic Ontologies and FIBO

by Oliver Browne, Nenad Krdzavac, Philip O'Reilly, Mark Hutchinson, David Saul, Dáire Lawlor, and Daragh McGetrick

Overview

Data, or the lack thereof, has become something of a hot topic in the realms of finance in recent years. Regulators have found that there is a significant gap between the information they require and the ability of companies to provide such data in a timely and accurate manner. This has led regulators to do what they inevitably do — regulate in a way that forces firms to provide more data than is actually required to reduce risk, principally in the form of BCBS (Basel Committee on Banking Supervision) 239.

The key principles of this regulation pose several problems for traditional financial institutions. The most glaring of these is the fact that the data simply may not be present in a format that can be quickly aggregated. Back-office reporting generally relies on legacy information systems that were really just short-term solutions to data problems and have since been wrapped or maintained rather than replaced. This use of dated systems also inhibits the integration of big data analytics and real-time data insight.[1] Redesigning old systems or developing new ones from the ground up seems to be where BCBS 239 points the sector. However, there are always alternatives to such a costly and time-consuming approach.

Ontologies have been explored as alternatives to traditional relational databases for some time, particularly in sectors such as manufacturing and microbiology. Their application in finance could possibly be game-changing due to several key benefits of an ontological approach:

- **Common language.** FIBO (Financial Industry Business Ontology) is a standard financial ontology language being developed to enable a common understanding among financial institutions. The goal of this dictionary of financial data is to establish a commonality among industry peers and financial products that would ultimately allow transactions between firms to take place free of human interaction, thereby forming the basis for automated "smart contracts." For regulators, it would also allow aggregation and comparison of data from all institutions on a like-for-like basis and enable them to fully grasp counterparty risk exposures.
- **Flexibility.** Data from core database systems is mapped to the ontology. This mapping can then be queried from multiple angles using a much more flexible approach than with traditional databases, also aggregating data from systems that would, without ontologies, be considered silos. New mappings can be added to extend existing data while retaining all existing data mappings. For example, if regulators require a new reporting template to be populated with data, rather than design a new one-off system with a view to completing this template, organizations can implement a new mapping that will populate it.
- **Domain knowledge.** Domain knowledge is key to the power of ontologies. Data that is seemingly unconnected or unknown in a traditional relational database (unless explicitly stated in the design process) can be inferred in an ontology approach. For example, say a legacy system has a blank field that should identify an instrument type. An ontology can assess the characteristics of the instrument and infer its type or use a class hierarchy to query at different levels of seniority of financial instruments.
- **Scalability.** As an ontology is a mapping of existing data, the required data can be streamed from database sources rather than having to occupy its own data warehouse. This streaming could potentially allow big data analytics over legacy system data. Overlaying ontologies onto "data lakes" makes them useful rather than just data dumps.

The What and Why of Ontologies

As we alluded to above, regulation has become one of the greatest demands on financial institutions since the economic crisis of 2007-2008 and the collapse of Lehman Brothers, with increased reporting and data requirements as well as stringent capital requirements. The introduction of Basel II, Basel III, Dodd-Frank, MiFID (Markets in Financial Instruments Directive) II, and the aforementioned BCBS 239 has significantly increased the strain on financial institutions' reporting capabilities. These increased reporting demands combined with a history of traditional financial institutions investing in back-office IT systems that focus on the short-term deliverable rather than a long-term solution have led to significant data quality issues, integration problems, high costs in maintaining legacy systems, and even regulatory fines. It is reported that, since 2009, failures in customer reporting have cost the world's top investment banks US $43 billion, making this the single most expensive compliance issue.[2]

For financial institutions, regulatory reporting has become something of a jigsaw puzzle — one that must be cobbled together into a coherent picture from several boxes into which the pieces from different puzzles have been put over time, for an audience that will never appreciate the pain involved in organizing that picture or the time and manpower required to build it. It can take a significant amount of time to complete reporting templates and even longer to verify their results — if that can be done at all. Regulators must then collect, collate, verify, and analyze the data received. Yet the data may not necessarily be "like for like" from one institution to another, and counterparty risk may not be fully transparent. This is where calls for a financial data standard have been growing and where Semantic Web technologies (ontologies) may be the answer.

The Semantic Web is the brainchild of Tim Berners-Lee,[3] who has proposed that data should all be connected by a linked data web in which relationships are displayed in a standard format and vocabulary to allow inference and to make data more meaningful and powerful.[4] An ontology is defined as "an engineering artefact, usually a model of (some aspect of) the world; it introduces vocabulary describing various aspects of the domain being modeled, and provides an explicit specification of the intended meaning of the vocabulary."[5] An ontology also allows us to infer relationships in a data set using reasoning, whereas databases cannot.[6] FIBO is a data standard language, currently in development, that aims to standardize the terms, relations, and language of the global financial industry. This standard is being proposed as a Rosetta Stone for financial data that will enable end-to-end processing of financial products.[7]

Legacy Systems and Reporting Requirements

Financial institutions have traditionally relied on short-term investments in information systems that aim to fulfill an immediate need rather than a strategic requirement, particularly for audits and regulatory reporting. This short-sighted strategy has led to various database systems that often contain data for which the documentation has been lost over time and whose designers have long since retired or left the company. Such legacy systems rely on a large amount of human input and verification of data, often producing a daily feed of physical paper printouts or spreadsheets, figures from which must be entered manually into a newer system that has been developed over time to fulfill a new task. This method of coping with new requirements is called "wrapping." The underlying legacy system is maintained, and outputs are fed into a newer system that can provide new capabilities. Something akin to this process can be seen in Figure 1.

It can take a significant amount of time to complete reporting templates and even longer to verify their results — if that can be done at all.

The idea of maintaining or redeveloping legacy systems can work, in theory, provided:

- Data requirements don't change.
- Regulations don't change.
- The underlying data is completely understood and documented.
- Business logic doesn't change.

In practice, rules and regulations are constantly subject to changes and updates, new calculations must be reported, and more detail must be provided to the regulator. This new required information often leads to large Excel-based macros that are designed to pull information from specific locations with standard file paths that must be updated monthly and will break if data is in the incorrect format, location, or sometimes even a different column.

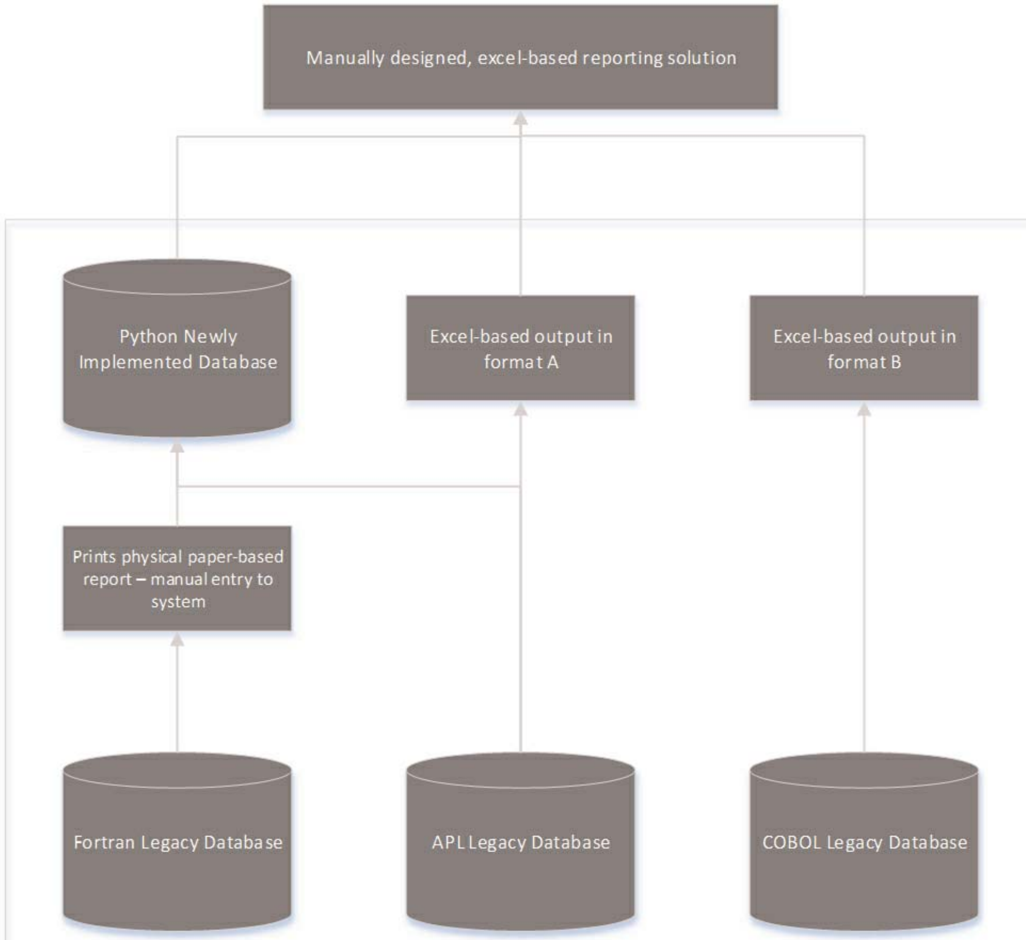

Figure 1 — Example of current regulatory reporting scenarios.

The dependency on siloed data and Excel-based manual aggregation and verification of data will cause many institutions to incur fines under BCBS 239, which mandates adherence to such principles as accuracy and integrity, completeness, timeliness, and adaptability of reported data.[8] Adaptability is one of the key concerns of institutions as current systems require a large amount of effort to complete ad hoc queries by regulators. Greater flexibility and a data standard are necessities for the future of financial reporting, for both institutions and regulators.

A Dictionary for Financial Data

As a proposed solution for the aggregation of data from separate sources, a process is being undertaken to develop a standardized language for financial data. FIBO is a dictionary and graphical representation of financial data that allows for the comparison of data on a like-for-like basis by using semantics. For example, let's say that two people are talking about equities, and one mentions an "ordinary share." In such a case, the other person can safely assume that this is semantically identical to a "common share." We infer this relationship based on our personal vocabulary; however, databases are unable to make this connection unless it is explicitly stated in the initial design of the database. Another example would be where a financial instrument has two occurrences in a data set but some details are different (e.g., the name has been manually entered incorrectly). A FIBO-enabled data structure would allow for the identification of errors in these two occurrences, whereas a traditional database would not, unless connected over other properties.

Database Architecture and FIBO

A key value-add of FIBO is the ability to leave existing IT architecture where it lies and use FIBO as an interpretive layer that sits on top of existing databases and aggregates data from multiple sources to provide output in a machine- and human-readable format (see Figure 2).[9] Businesses are generally quick to patch up

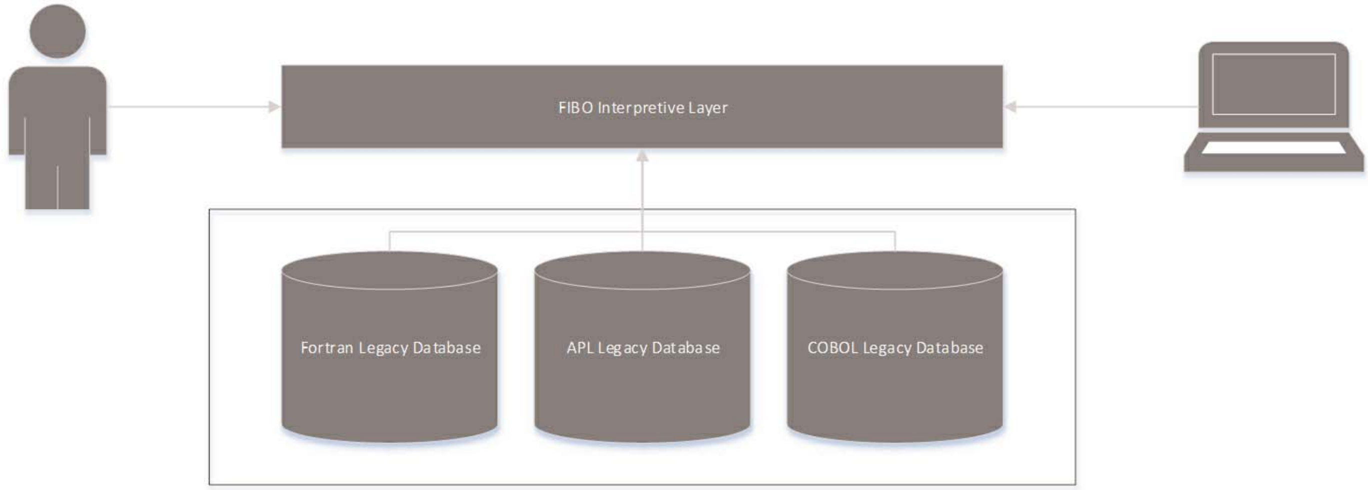

Figure 2 — User and system query of databases using FIBO.

old systems, wrap them with new ones, or even completely redevelop new systems. However, these efforts are generally undertaken without a complete understanding of the underlying data or sufficient documentation of the system. Maintaining old systems can be a costly process; however, the development of a new system from scratch is a massive undertaking that can cost large sums of money and, more importantly, time.[10] Once a new database system is designed and implemented, changing any business logic or rules within that system can become a daunting task, as changing the existing and implemented schema is extremely complex and error-prone if original documentation has not been maintained or the designers of the system have left the organization.

Flexibility and Mapping

FIBO acts as an interpretive layer of mapping over financial data and provides significant flexibility in the implementation of the standard. Key to the process of implementation is the mapping of meanings of data from systems to FIBO. This is a complex process that requires intimate knowledge of the technology involved in the use of the ontologies, the systems and data being mapped, and FIBO itself. However, once mapping has been completed a single time for any instance of data, then that is all that is required to use this data and apply the same mapping to semantically identical instances. Mapping is also flexible regarding regulations and reporting of data. For example, say a financial instrument that falls into a certain category for reporting is changed in regulations. Rather than having to manipulate this in Excel, update macros, or reengineer the business logic and schema, the mapping of this instrument can simply be updated to feed elsewhere. This redirection can be done by business analysts without resorting to specialized database management skills.

Ontologies are also flexible in the use of characteristics and properties in mapping data. The use of class hierarchies allows for a more complete data structure. Hierarchies allow financial instruments to adopt characteristics of the classes that supersede them as well as being able to adopt characteristics that are unique to themselves. For example, Figure 3 displays the relationships and properties in a class hierarchy of two different types of shares, a common share and a preference share. Both types of shares have the characteristics of the Share class but also have their own unique properties that separate them. The power of this representation is the ability to query using inference, to query both by individual instrument type or higher up the hierarchy if necessary. A traditional database may not relate a common share and a preference share as types of the Share class unless this is explicitly stated. An ontology, in contrast, can infer this relationship. Also shown in Figure 3 is the relationship between preference shares and ordinary shares, the former being senior to the latter. Semantically, we can infer that common shares are junior to preference shares, but a database cannot infer this relationship unless it is explicitly stated.

Another benefit of ontologies is the ability to apply flexible queries over databases that may not normally be interfaced. Such integration issues generally mean that data dumps must be taken from multiple sources and collated using Excel-based methods. A 2016 Deloitte report found that compliance teams spend too much time collecting data and not enough time analyzing it.[11] FIBO allows flexible querying over SPARQL, an industry standard query language.

Visual Representation

An ontology provides a graphical representation of data in the form of RDF (Resource Description Framework),[12] a standard for representing graphical data that produces linked, labeled, directed multi-graph triples. A triple is a subject, predicate, and object relationship between data. In financial terms, an example would be: a fixed coupon bond is issued by a sovereign government (see Figure 4). These triples form an extensive web or graph of relationships and attributes that ultimately represent a previously unseen view of data and its connectivity that can be searched both visually and over query language. These graphical representations lend themselves to visualization in graph hardware processors, offering enhanced performance and analysis through visualization.

FIBO breaks financial instruments and their relationships down to the most basic level of detail available. In most instances, this base level is a contract between two parties, and relationships and details are built from this point. There are two key principles of this method: the first is to enable comparability between instruments from a base level, and the second is to allow for the most accurate measure of counterparty risk, something that recent regulator investigations have exposed as being dangerously deficient in detail.[13] BCBS 239

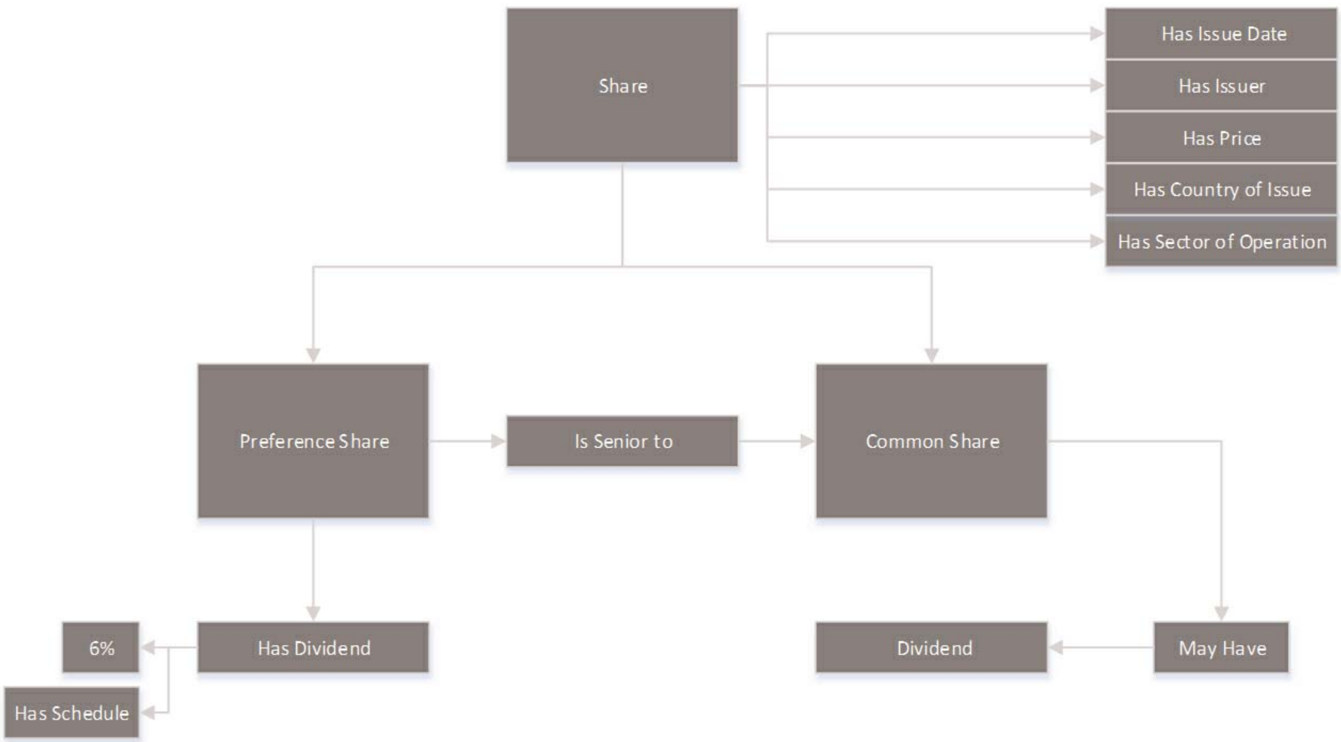

Figure 3 — Share class hierarchy in FIBO.

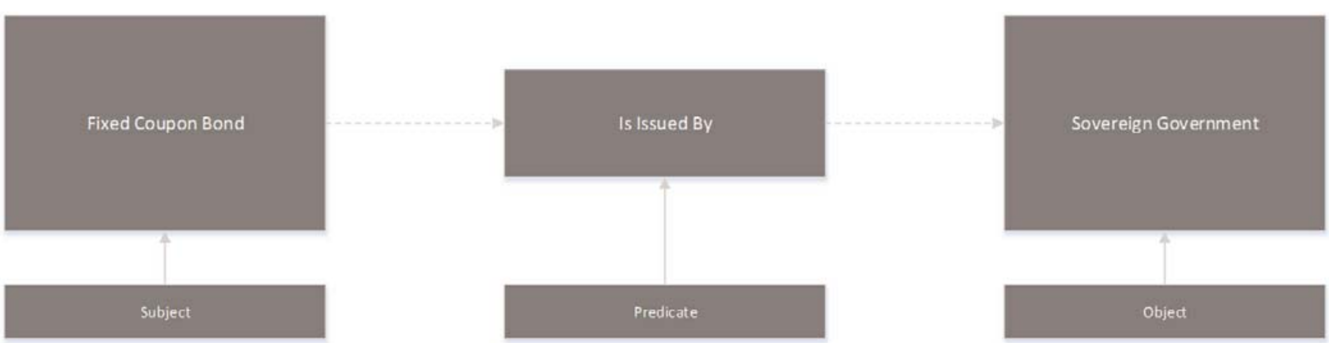

Figure 4 — RDF triple representation of relationships.

requires a vast amount of data to be aggregated and the minutiae of counterparty risk exposures to be provided and reported.

Visualization is one of the key benefits of adopting an ontological approach to data representation. Imagine an Excel data sheet that contains 250 columns and 25,000 lines; this sheet contains 6.25 million data points that an analyst must assess and investigate. Representing this vast amount of data in tabular format doesn't allow for in-depth understanding of the detail and intricacies of the relationships involved, causing a large amount of time to be spent reformatting the data from lines and columns to a more agreeable and understandable format. Ontologies simplify this process through visualization of data, which is key to understanding complex webs of big data. A simplified ontological view of an investment fund can be seen in Figure 5.

FIBO State of Art

FIBO is rapidly gaining momentum. As a community-driven standard being developed by the EDMC (Enterprise Data Management Council) on behalf of the OMG (Object Management Group), FIBO relies on testing and feedback to develop the standard further. OMG has a proven track record in delivering ontology standards with the production of UML2 (Universal Modeling Language 2), among other leading standards.

A number of organizations have completed multiple semantic demonstrations, proofs of concept, and practical implementations of FIBO. State Street is involved in the application and development of this technology, with an application in interest rate swaps[14] and most recently in bonds and equities. Figure 6 shows the taxonomy (class hierarchy) of implemented instruments in the most recent bonds and equities implementation. Extensions, discoveries, and developments from projects such as this are fed back to OMG and the EDMC to continuously improve and develop the standard for industry-wide adoption.

Conclusion

Financial institutions face a new world order of mass regulation and ever more demanding queries from regulators, which strain legacy systems and require a large amount of time and effort to piece together. For their part, regulators currently have to work to homogenize the data they receive in terms of formatting and aggregation. Given these challenges as well as the requirements of BCBS 239, there is a significant demand from regulators to be supplied with homogenous data and from financial institutions for an alternative to the development of data lakes and warehouses and IT architecture redesign.

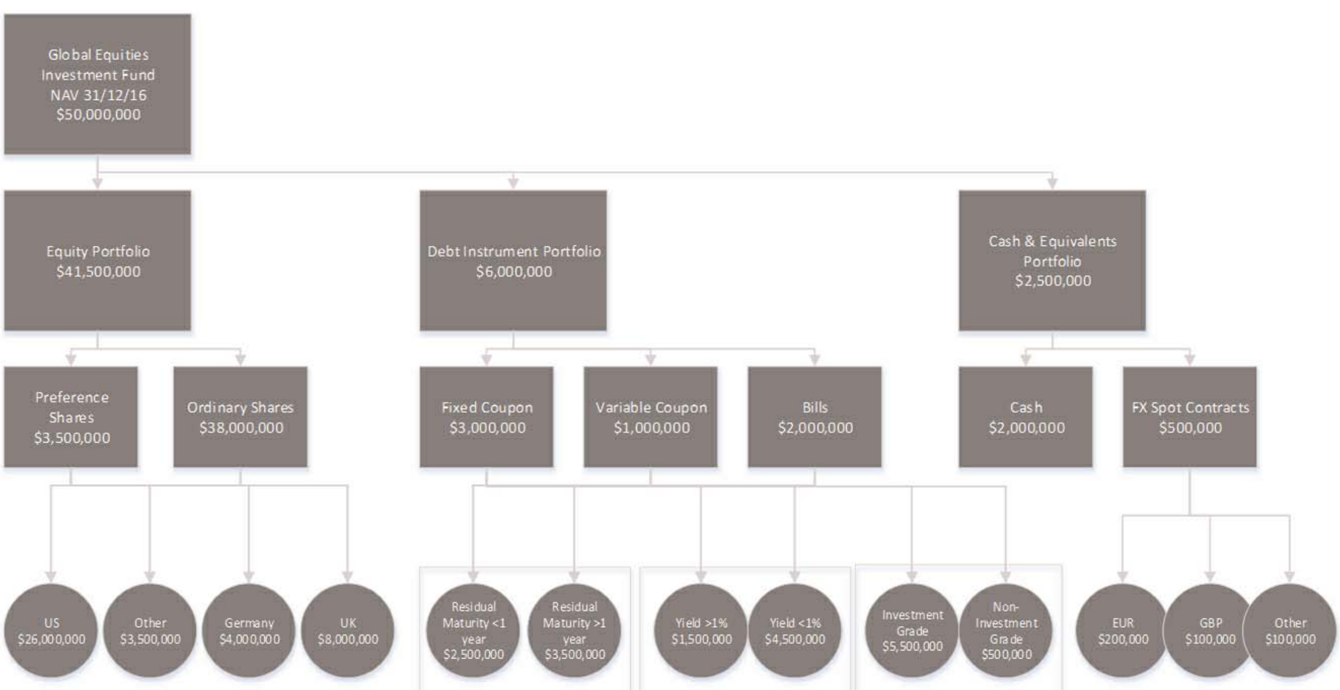

Figure 5 — Investment fund example in FIBO graph format.

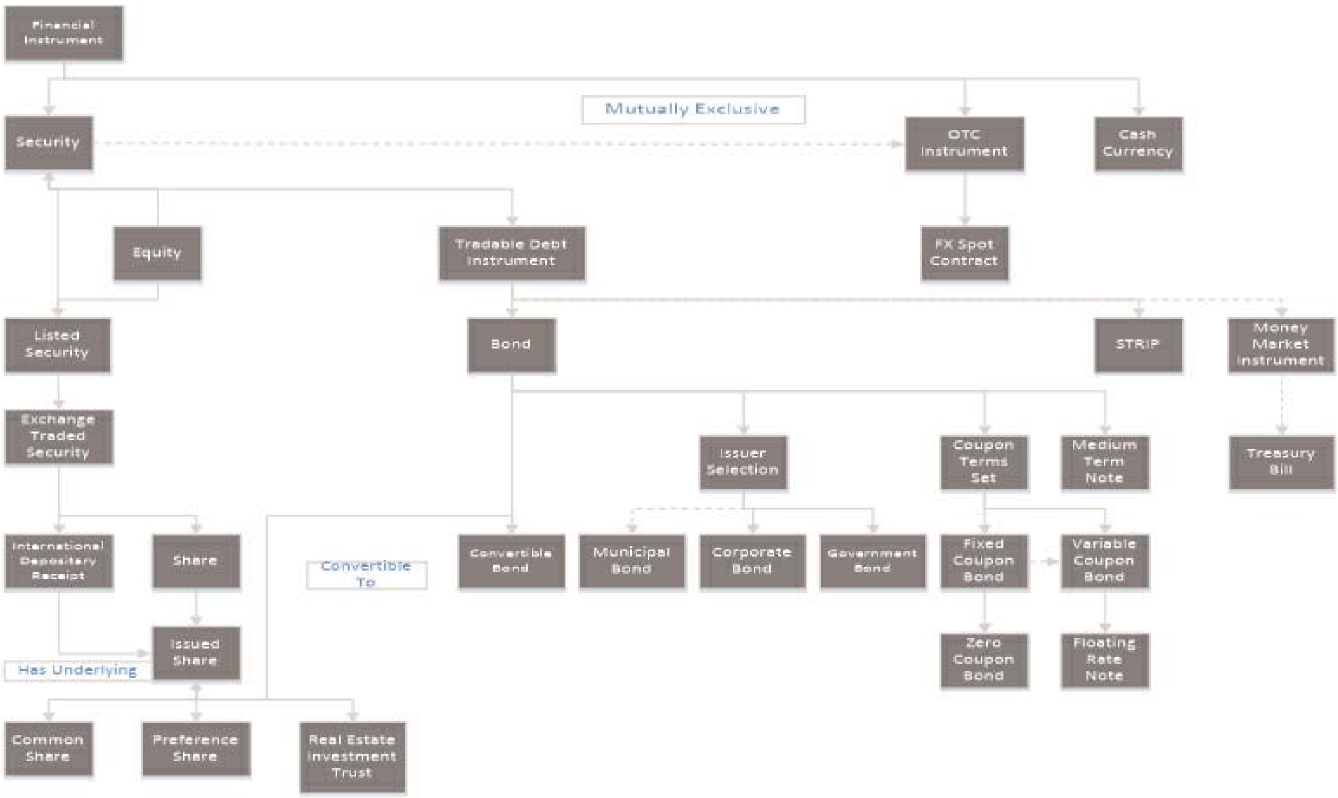

Figure 6 — Implemented FIBO taxonomy in a State Street proof of concept.

FIBO is an option with great promise to solve the coming headaches with a flexible and scalable technology that will last long into the future and be as malleable as the regulatory environment requires. Rather than designing a complete new system, schema, and architecture for a narrow purpose, organizations can implement new mappings. Queries can be answered in real time using a standard query language rather than data dumps and Excel-based methods. Complex data can be visualized to allow greater, deeper, and faster understanding and insights. Most importantly, historically siloed systems and data can be integrated in a universal data language, allowing much greater control over and insight into data. FIBO is still in development, but the implementation of the technology is well established in other industries. Its application and adoption in finance will be extremely beneficial to all parties.

Endnotes

[1]Shah, Nimish. "Legacy Systems — Still the Big Barrier to Banking Innovation." *Banking Technology*, 14 December 2015 (www.bankingtech.com/412822/legacy-systems-still-the-big-barrier-to-banking-innovation/).

[2]Noonan, Laura. "Customer Reporting Failures Top Big Bank Fines." *Financial Times*, 27 March 2016 (www.ft.com/content/7993ab16-f2c0-11e5-aff5-19b4e253664a).

[3]"Vocabularies." W3C, Semantic Web (www.w3.org/standards/semanticweb/ontology).

[4]Berners-Lee, Tim. "The Next Web." TED Talks, February 2009 (www.ted.com/talks/tim_berners_lee_on_the_next_web).

[5]Horrocks, Ian. "What Are Ontologies Good For?" Chap. 9 in *Evolution of Semantic Systems*, edited by Bernd-Olaf Küppers, Udo Hahn, and Stefan Artmann. Springer-Verlag Berlin Heidelberg, 2013 (www.springer.com/us/book/9783642349966).

[6]Horrocks, Ian. "Ontologies and the Semantic Web." *Communications of the ACM*, Vol. 51, No. 12, December 2008 (http://dl.acm.org/citation.cfm?doid=1409360.1409377).

[7]Crosman, Penny. "State Street Tests a 'Rosetta Stone' for Bank Databases." *American Banker*, 14 July 2016 (www.americanbanker.com/news/state-street-tests-a-rosetta-stone-for-bank-databases).

[8]"BCBS 239: Risk Data Aggregation and Reporting." EY, February 2015 (www.ey.com/Publication/vwLUAssets/EY-bcbs-239-risk-data-aggregation-reporting-AU/%24FILE/EY-bcbs-239-risk-data-aggregation-reporting-AU.pdf).

[9]Fortineau, Virginie, Thomas Paviot, and Samir Lamouri. "Improving the Interoperability of Industrial Information Systems with Description Logic-Based Models — The State of the Art." *Computers in Industry*, Vol. 64, No. 4, May 2013 (https://doi.org/10.1016/j.compind.2013.01.001).

[10]Hainaut, Jean-Luc, et al. "Migration of Legacy Information Systems." Chap. 6 in *Software Evolution*, edited by Tom Mens and Serge Demeyer. Springer-Verlag Berlin Heidelberg, 2008 (https://doi.org/10.1007/978-3-540-76440-3_6).

[11]"Forward Look: Top Regulatory Trends for 2016 in Banking." Deloitte, 2016 (www2.deloitte.com/content/dam/Deloitte/us/Documents/finance/us-outlooks-top-regulatory-trends-for-2016-in-banking-reg.PDF).

[12]W3C (see 3).

[13]"Where Have You Been All My Life? How the Financial Services Industry Can Unlock the Value of Big Data." PwC FS Viewpoint, October 2013 (www.pwc.com/us/en/financial-services/publications/viewpoints/unlocking-big-data-value.html).

[14]David, Emilia. "FIBO Marches Forward: A Look Inside State Street's FIBO Proof of Concept." Waters Technology, 3 June 2016 (www.waterstechnology.com/waters/feature/2459451/fibo-marches-forward-a-look-inside-state-streets-fibo-proof-of-concept).

Oliver Browne is a PhD researcher in accounting at Cork University Business School, University College Cork (UCC). He attended UCC from 2009 to 2014 and achieved a bachelor's degree in finance and a master's degree in corporate finance. He has worked in the banking sector with Bank of Ireland, Group Treasury, in the preparation and submission of regulatory reporting and financial statements and the management of liquidity in group deposits. He previously worked with Allied Irish Banks, Corporate Institutional and Commercial Banking, in the area of loan book management and risk reporting. In 2016, Mr. Browne joined UCC as a researcher in regulatory and risk reporting. He can be reached at oliver.browne@ucc.ie.

Nenad Krdzavac is an accomplished research scientist with extensive experience conducting ground-breaking research, contributing to knowledge of Semantic Web technologies. Dr. Krdzavac has strong research interests spanning description logics, automated reasoning, ontology development, Web services development, and model-driven engineering. He previously held roles at Michigan State University and the University of Belgrade. Dr. Krdzavac is a co-inventor of a methodology and system for inferencing semantic equivalence between concepts (US patent pending). He can be reached at Nenad.Krdzavac@ucc.ie.

Mark Hutchinson is Professor of Finance and Deputy Head of the Cork University Business School at University College Cork (UCC). From 2000 to 2002, Dr. Hutchinson traded futures, options, and cash equities at HSBC Investment Bank in London. At HSBC he was responsible for institutional market making and proprietary trading. In 2002 he joined UCC, where he has led several research teams, securing over €5 million in research funding. With researchers at Cass Business School (City University of London), Dublin City University, UCC, and the University of Miami, he founded the Centre for Investment Research, which focuses on researching investments and trading strategies. More recently he has led several fintech research projects with industry. He regularly presents his research throughout Europe and the US and has published in leading international academic and practitioner journals. He currently sits on the curriculum development committee of the CAIA (Chartered Alternative Investment Analyst) Association in the US. He can be reached at m.hutchinson@ucc.ie.

Philip O'Reilly is a Senior Consultant with Cutter Consortium's Business Technology & Digital Transformation Strategies and Data Analytics & Digital Technologies practices. He is also Director of the State Street Advanced Technology Centre at University College Cork (UCC), the Academic Co-Director of Boole Business Labs, and a Senior Lecturer in Business Information Systems at UCC. Dr. O'Reilly received the Stafford Beer Medal in recognition of the most outstanding contribution to the philosophy, theory, and practice of information systems from the Operational Research (OR) Society at an awards ceremony that took place at the Royal Society in London. He has been invited to deliver keynotes and workshops by numerous multinational companies and at leading practitioner events, including the Banking & Payments Federation of Ireland National Conference. *Dr. O'Reilly has been a leading member of research teams that have been successful in securing in excess of €2 million in research funding in recent years. His research concentrates on digital business, with a specific focus on financial services. Dr. O'Reilly's work has been published in leading journals in the information systems field, including* European Journal of Information Systems, Journal of Strategic Information Systems, *and* Information Technology and People. *He can be reached at poreilly@cutter.com.*

David Saul is Senior VP and Chief Scientist at State Street Corporation, reporting to the CIO. In this role, he is responsible for proposing and assessing new advanced technologies for the organization as well as evaluating technologies already in use at State Street and their likely evolution in order to reinforce the organization's leadership position in financial services. Mr. Saul previously was Chief Information Security Officer, in which he oversaw State Street's corporate information security program, controls, and technology. Prior to that, he managed State Street's Office of Architecture, where he was responsible for the overall enterprise technology, data, and security architecture of the corporation. Mr. Saul joined State Street in 1992 after 15 years with IBM's Cambridge Scientific Center, where he managed innovations in operating systems virtualization, multiprocessing, networking, and personal computers. He serves as a Trustee of the Massachusetts Eye and Ear Infirmary. In 2007 Mr. Saul was honored with a Computerworld *Premier 100 IT Leader Award, and* American Banker *named him one of its Top Innovators of 2013 for his work on data semantics and standards. Mr. Saul holds bachelor's and master's degrees from MIT. He can be reached at dnsaul@statestreet.com.*

Dáire Lawlor, who joined State Street in 2010, draws on over 25 years of management experience working on both the client and the vendor side across a range of sectors in Ireland, the UK, Spain, and India. With over a decade's worth of experience in senior management roles in the asset and fund management sector, he is currently head of the Consultancy Services Group for State Street International (Ireland) Ltd., managing a team of expert technicians, analysts, and project managers whose remit includes innovation, fintech/regtech, and data solutions programs as well as a variety of new business, product development, client change, and internal transformation initiatives. He can be reached at DLawlor@statestreet.com.

Daragh McGetrick works for the Consulting Services Group in State Street Global Services. He has 13 years' experience in financial services, with the last 10 years focused on transformation initiatives across fund accounting, derivatives pricing, custody, financial reporting, and regulatory reporting. Mr. McGetrick holds a bachelor's degree in economics and finance, a higher diploma in IT, and a master's degree in economic science. He can be reached DMcGetrick@statestreet.com.

How Robo-Advisors Manage Investment Portfolios

by Jie Yang, Hanxi Ye, Yadan Wei, and Linqian Bao

What Is a Robo-Advisor?

Robo-advisors are a class of financial advisors that provide financial advice or portfolio management online by employing algorithms such as modern portfolio theory (MPT) to conduct portfolio management.[1] Not that automated, online investing platforms are a new thing. In 1996, Nobel Prize winner William F. Sharpe co-founded Financial Engines, one of the first algorithmically driven investment platforms. Now, however, robo-advisory (digital investment) platforms are becoming more visible. Wealthfront, Betterment, and traditional corporations such as Charles Schwab, Vanguard, and Fidelity have also launched their own robo-advisors.[2] In fact, more than 50 firms can loosely be included in the robo-advisor space.

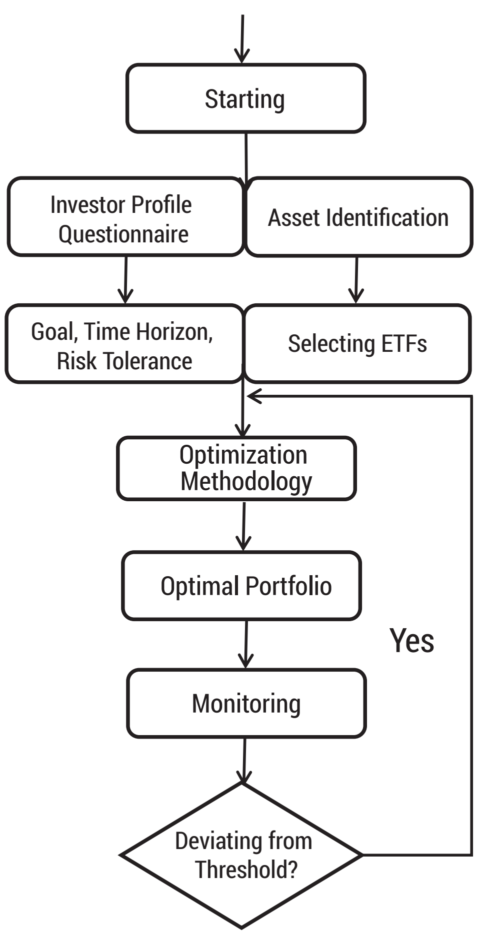

Figure 1 — The asset allocation process in robo-advisors.

In 2015, robo-advisors managed US $55-$60 billion in assets[3] and Schwab Intelligent Portfolios, Wealthfront, and Betterment make up more than 50% of this robo-advisory services market.[4] In the first quarter of 2016, Financial Engines had surpassed $975 billion in assets under contract and $122 billion in assets under management for over 9.6 million people.[5] According to a survey conducted by A.T. Kearney, robo-advisors will run $2 trillion assets under management by 2020.[6]

In general, robo-advisors are classified into three types:[7]

1. Pure technology websites that are devoid of human advisors
2. Companies that employ advisors to use technology only to communicate with clients
3. Traditional financial services companies that have recently expanded their online advice offering

Robo-advisors differ from each other in provided services, fees, and (especially) algorithms, including asset classification, exchange-traded fund (ETF) selection, portfolio optimization, and rebalancing.

Investment Methodology

The investment methodology of most robo-advisors begins with asset allocation, proceeds to implementation, followed by monitoring and rebalancing. The flowchart in Figure 1 illustrates the stages involved in a robo-advisor system.

Below, we choose the most notable robo-advisors in the US — namely, Betterment, Wealthfront, and Schwab Intelligent Portfolios — to make detailed comparisons. Although they are quite similar, they have their own features. Betterment uses smart financial models, such as a downside-risk optimization model, to make sure that each of its four primary investing goals (retirement, safety net, general investing, and major purchase) is based on a different stock allocation path. Wealthfront is more like a technology company, making sophisticated financial advice and investment strategies available to every individual investor at low cost by combining financial theory and machine learning. Schwab, as a traditional financial services firm, uses more traditional investment strategies in asset allocation.

Asset Class Identification

Given the different features of asset classes, robo-advisor companies often divide them into three broad categories: stocks, bonds, and inflation assets. Stocks have high expected returns in the long run but high volatilities, while bonds have low expected returns but low volatilities, as well as low correlation with stocks. We also need assets that are highly correlated with inflation, such as US Treasury Inflation-Protected Securities (TIPS), real estate, and natural resources.

Because of increasing external shocks, lower returns, and higher correlations between traditional assets, Schwab's asset allocation also encompasses nontraditional asset classes, like gold or other commodities. All the asset classes in Schwab Intelligent Portfolios must be accessible through at least two liquid ETFs and be minimally correlated to achieve greater diversification benefits. Asset classes in Schwab are divided by their role in the portfolio. According to Schwab:[8]

- Growth potential will come primarily from the equity allocations.
- Growth potential and income will come from dividend-paying stocks.
- Income will come from a broad array of fixed-income investments.
- Inflation protection comes from allocations to US inflation-protected bonds, FDIC-insured cash, and Real Estate Investment Trusts (REITs).
- Defensive assets are generally assets that have low or negative correlations with equity securities.

When building the portfolio, Schwab combines market-cap strategies and fundamental strategy by dividing stocks into market-cap stock and fundamental stock. For example, we can divide US large company stocks into US large company stocks and US large company fundamental stocks. For its part, Betterment classifies assets into stocks and bonds in its portfolios, while Wealthfront's portfolios include real estate and natural resources besides stocks and bonds.

In general, the main goal of asset classification is to construct a diversified portfolio; thus, the main methodology behind this classification is to estimate the correlation between the asset classes. From Table 1, we can see that Schwab has the largest number of asset classes and is also the only robo-advisor invested in master limited partnerships. However, there is always a tradeoff between diversification and expected return. By including some lower-risk assets (e.g., bonds) in the portfolio, you may be losing out on the potentially higher return of assets such as stocks. And with too many asset classes, using the mean-variance model to estimate the optimal portfolio often brings many problems due to the maximum and minimum constraints, as we will discuss below.

The main goal of asset classification is to construct a diversified portfolio; thus, the main methodology behind this classification is to estimate the correlation between the asset classes.

Selecting ETFs

Mutual funds can make no claim to be superior over the market averages.[9] Most robo-advisor companies exclusively use open-ended index tracking ETFs rather than mutual or closed-end funds due to their low manager risk, low embedded costs, and natural tax efficiencies. They select ETFs based on two main features: cost and liquidity. The cost of ETFs consists of fund expenses imposed by the ETF administrator, bid-ask spread in market interaction, and performance deviations from the benchmark index, or tracking error. Sufficient liquidity is needed to make sure the selected ETFs can be withdrawn at any time.

When assessing whether an ETF accurately reflects one asset class, Schwab considers geographic exposure, market cap, industry/sector concentration, yield, duration, and tracking error. At the present time, exchange-traded notes (ETNs) and limited partnership ETFs are ineligible for Schwab Intelligent Portfolios. Once ETFs are selected for Schwab Intelligent Portfolios, they are monitored on a periodic basis to ensure that they remain an appropriate choice for the portfolios.[10]

Betterment's approach is to aggregate expense ratio, bid-ask spread, and tracking error into a single variable called total annual cost of ownership:[11]

$$Total\ Cost = \omega_E * E + \omega_S * S + \omega_T * T$$

<table>
<tr><th colspan="2">Schwab</th><th colspan="2">Betterment</th><th colspan="2">Wealthfront</th></tr>
<tr><td rowspan="2">Growth</td><td rowspan="2">• US large company stocks
• US small company stocks
• International/ developed country/ large company stocks
• International/ developed country/ small company stocks
• International emerging market stocks
• US large company stocks (high dividend)</td><td rowspan="2">Stocks</td><td rowspan="2">• US large-cap value stocks
• US mid-cap value stocks
• US small-cap value stocks
• International/ developed country stocks
• Emerging market stocks</td><td>Stocks</td><td>• US stocks
• Foreign-developed stocks
• Emerging market stocks</td></tr>
<tr><td rowspan="3">Fixed Income</td><td rowspan="3">• Dividend stocks
• Municipal bonds
• Corporate bonds
• Emerging market bonds</td></tr>
<tr><td>Growth and Income</td><td>• International/ developed country/ large company stocks (high dividend)
• Master limited partnerships (MLPs)
• US investment-grade corporate bonds</td><td rowspan="4">Bonds</td><td rowspan="4">• Short-term treasuries
• Inflation-protected bonds
• US municipal bonds
• US investment-grade bonds
• International/ developed country bonds
• Emerging market bonds
• US high-quality bonds</td></tr>
<tr><td>Income</td><td>• US corporate high-yield bonds
• US securitized bonds
• International emerging market bonds
• Preferred stocks
• Bank loans and other floating-rate notes
• US inflation-protected bonds</td></tr>
<tr><td>Inflation</td><td>• US REITs
• International REITs
• Treasuries</td><td>Real Estate</td><td></td></tr>
<tr><td>Defensive Assets</td><td>• Gold and other precious metals
• International/ developed country bonds</td><td>Natural Resources</td><td></td></tr>
</table>

Table 1 — Robo-advisors: asset classification.

Here, E denotes expense ratio, imposed by the ETF administrator, S denotes bid-ask spread, and T denotes the tracking error. Once this variable has been calculated, Betterment then selects the ETFs with the least cost.

Wealthfront employs the same methodology, selecting passive index ETFs that either minimize lending or share the lending revenue with their investors to reduce management fees.

Investor Profile

In order to match each client's investment goals and determine the optimal mix of the portfolio, robo-advisor companies often ask clients risk questions and verify the consistency among the answers. A typical survey consists of dozens of questions to identify an individuals' risk tolerance, while some companies, such as Wealthfront and Schwab, have simplified this process (see Appendix).

Investment Goal and Portfolio Sets

In Schwab's questionnaire, investors can choose from among these goals:

1. Prepare for retirement
2. Save for major upcoming expenses (education, healthcare bills, etc.)
3. Save for something special (vacation, new car, etc.)
4. Build a rainy day fund for emergencies
5. Generate income for expenses
6. Build long-term wealth

These six goals have different portfolio sets. In general, the portfolio sets of goals 1, 3, and 6 all include stocks, fixed income, cash, and commodities, while only stocks, fixed income, and cash are included in the portfolio sets of goals 2, 4, and 5.

Betterment is a goal-based investing company that claims to customize the time horizon and purpose for every investment account to optimize investors' savings and ensure that investors meet their goals. To manage risk for individual goals, Betterment has developed four primary types of investing goals: retirement, safety net, general investing, and major purchase. Each of these goal types is based on a different stock allocation glide path, which provides an allocation plan based on the investor's personal time horizon for each goal. Betterment divides its asset classes only into stocks and fixed income (six stocks and six bonds). So for each of the goal types, Betterment provides a recommended minimum and maximum stock allocation, anticipated term, and certain cash-out assumptions.[12] Table 2 shows the detailed goals, which we will discuss further below.

Wealthfront just divides the investor's goals into general savings, retirement, and other. Since Wealthfront enforces minimum and maximum allocation constraints for each asset class, the portfolio sets have almost the same asset classes, while the weight of each asset class in the specific portfolio set differs. That said, taxable accounts always include municipal bonds, while retirement accounts replace municipal bonds with corporate bonds and emerging market bonds.

Risk Spectrum

Educational psychologist Michael Roszkowski and his coauthors found that combining risk-attitude (or behavior-related) questions with objective questions will provide a more complete understanding of the investor.[13] Schwab thinks the investor's risk capacity is based on information about their financial situation, while their willingness to accept risk is typically indicated by the level of volatility they're comfortable with and other factors. And an investor's risk capacity and risk willingness are generally independent of each other:

Type of Goal	Stock Allocation Range	Time Horizon
Retirement	(56%, 90%)	Up to 50 years
Retirement Income	(30%, 56%)	Up to 30 years
Safety Net	40%	Rolling term
General Investing	(55%, 90%)	Indefinite
Major Purchase	(5%, 90%)	Up to 30 years

Table 2 — Betterment portfolio sets according to investment goal.

Risk Tolerance = Willingness to Take Risk + Capacity to Bear Risk

To assess an investor's attitudes toward risk, Schwab's Intelligent Portfolios IPQ (Investor Profile Questionnaire) contains 15 questions:[14]

- Five questions capture factual information to assess an investor's risk capacity.
- Six questions collect information about the investor's behavioral attitude toward risk to calculate their risk willingness score.
- Four questions ask about the investor's age and product preferences, such as Total Return/Income or Muni/Taxable bonds.... [to] help determine the appropriate set of portfolios.

Behavioral economics research shows that individuals consistently overstate their true risk tolerance — especially male investors who are educated and overconfident.

Betterment doesn't measure investors' risk tolerance directly. Instead, it tries to do asset liability management, which is a framework for ensuring that the investor's future expenditures are funded (if not overfunded) whenever they appear. Betterment only asks its investors four questions to understand their goal and time horizon. It then uses a downside risk model to build recommendations.

In order to identify the ideal asset allocation for an investor's needs, Wealthfront tries to pinpoint an investor's risk tolerance with eight questions. It also tries to calculate the objective capacity to take risk and subjective willingness to take risk independently, just as Schwab does. However, unlike Schwab's methodology of assigning equal weight to "Risk Willingness" and "Risk Capacity," Wealthfront assigns a heavier weight to whichever component is more risk-averse, because behavioral economics research shows that individuals consistently overstate their true risk tolerance — especially male investors who are educated and overconfident.[15] Clients are also allowed to adjust their assigned risk score once every 30 days.

Asset Allocation Models

Among the complicated asset allocation models used by robo-advisor companies, the mean-variance model and Black-Litterman model are the most common. Mean-variance optimization, first introduced by economist Harry Markowitz in 1952, tries to maximize the expected return for a given level of portfolio risk. And the Black-Litterman model, developed by economists Fischer Black and Robert Litterman at Goldman Sachs in 1990, combines managers' views about asset class returns with the market equilibrium implied returns, which are the set of returns that clear the market. We can derive the equilibrium returns from the formula below:

$$\Pi = \delta \Sigma \omega_{\mu\kappa\tau}$$

where:

- Π is the vector of implied equilibrium excess returns for each asset.
- δ is the risk aversion coefficient.
- $\omega_{\mu\kappa\tau}$ is the market capitalization weight.
- Σ is the covariance matrix of the excess returns for the assets. A presumed efficient benchmark is included in the calculation.

As for these three robo-advisors, Schwab combines mean-variance optimization and full-scale optimization together to model the investor's preference as to loss aversion. The two techniques Betterment uses are the Black-Litterman model and the downside risk optimization model. In addition to the mean-variance optimization model, Wealthfront uses such techniques as the capital asset pricing model (CAPM) and the Gordon Growth model.

Mean-Variance (MV) Model

There are four major assumptions for Markowitz's MPT:

1. There are n risky assets in the market, with known distribution of their rate of returns, and a non-singular covariance matrix.
2. There are no frictions in the market. That is to say, we ignore any kinds of transaction costs, assume no bid-ask spreads, allow short-selling, and so on.
3. Individuals in the market make their decisions according to the mean and variance of financial assets. Specifically, they are risk-averse in that they prefer a portfolio with lower variance for a given level of risk, or they want to minimize risk for a given level of returns.
4. All individuals have the same time horizon for investment. So our analysis is based on a one-period, discrete time situation.

For each asset i, we denote its rate of return a random variable r_i, and define the random vector $\mathbf{r} = (r_1, r_2, ..., r_n)^T$. Let $\mathbf{e} = E(r)$ and $\Sigma = \text{Cov}(\mathbf{r})$ be its mean and covariance.

Then, we define vector $\omega = (\omega_1, \omega_2, ..., \omega_n)^T$ as the set of weights associated with a portfolio.

From our assumptions, we can formulate the optimal portfolio selection problem as follows:

$$\min_{\omega} \frac{1}{2} \text{Var}\{\omega^T \mathbf{r}\} = \frac{1}{2} \omega^T \Sigma \, \omega$$
$$\text{subject to } \omega^T e = E(\boldsymbol{r}_p), \omega^T I = 1$$

where $I = 1, 1, ..., 1)^T$ is the vector of ones.

We can solve the optimization problem via the method of Lagrange multipliers.[16] Firstly, the formulas can be rewritten as below using Lagrange multipliers:

$$\min \left\{\frac{1}{2} \omega^T \Sigma \omega + \gamma_1 [\, E(\boldsymbol{r}_p) - \omega^T e \,] + \gamma_2 (\, 1 - \omega^T I \,) \right\}$$

where γ_1 and γ_2 are the multipliers.

Then we can compute the three first-order conditions (FOCs) for a critical point as to ω, γ_1, and γ_2. Last, the optimal weight can be obtained with the three FOC equations. MIT finance professor Robert C. Merton has provided detailed steps of the analytical solutions.[17]

While the mean-variance model is a classic framework for asset allocation, Black and Litterman have pointed out that the use of it by investment managers has been limited by two recognized problems. First, it's not easy to estimate the expected returns of all assets and currencies for investors, with the result that their views are augmented. Second, the weights are extremely sensitive to the return assumptions used. For example, the optimization portfolio always ordains large short positions when investors impose no constraints. And the optimization portfolio often prescribes "corner" solutions with zero weights in many assets when short positions are ruled out.[18]

Capital Asset Pricing Model (CAPM)

In the early to mid-1960s, the famous capital asset pricing model was introduced, independently, by investment theorist Jack Treynor, Stanford finance professor William F. Sharpe, Harvard Business School professor John Lintner, and Norwegian economist Jan Mossin to explore the relation between risk and returns of assets under market equilibrium. Compared to the mean-variance model, it has two additional assumptions:

1. There is a risk-free asset in the market with no bid-ask spread.
2. All investors have access to the same information and agree about the risk and expected return of all assets (homogeneous expectations assumption).

The risk of a portfolio comprises systematic risk and unsystematic risk. Unsystematic risk can be diversified away to zero while systematic risk is undiversifiable. The critical improvement of CAPM is that it studies the pricing of a portfolio from the market perspective. CAPM introduces the concept of a market portfolio to reflect the investment choice of the whole market. Then, it uses the corollary that the market portfolio is optimized on the efficient frontier. That is to say, the market portfolio contains only systematic risk, or the total risk of the market portfolio can be viewed as systematic risk (beta, β). Then the CAPM prices portfolios via the expected market excess return and the portfolio systematic risk (beta, β). We denote the return of the market portfolio as r_m and risk-free rate as r_f Then we can have expected market excess return $E(\mathbf{r}_m) - r_f$. For any portfolio with return r_q, CAPM gives its pricing formula as follows:

$$\text{E}(\mathbf{r}_q) = r_f + \beta_{qm}(\text{E}(\mathbf{r}_m) - r_f)$$

where:

$$\beta_{qm} = \frac{Cov(\mathbf{r}_q, \mathbf{r}_m)}{Var(\mathbf{r}_m)}$$

Therefore, if the market portfolio is observed (i.e., we know the mean and variance of the market portfolio), we can calculate a portfolio's expected return by only examining the covariance between this portfolio and the market portfolio.

Although the CAPM offers powerful and intuitively pleasing predictions about how to measure risk and the relation between expected return and risk, University of Chicago finance professor Eugene F. Fama has argued that the empirical work shows that the relation between beta and average return is flatter than predicted by the CAPM.[19]

To improve the model, relaxing the homogeneous expectations assumption might be a good idea. In fact, an investment manager gets a lot of views about returns of assets from their analysts. Therefore, it is natural to find a way to incorporate subjective views about the outlook for global assets with the market equilibrium implied returns to obtain an optimal allocation. We will discuss this in more detail in the following section.

Black-Litterman (BL) Model

Now, let μ, Σ, δ be the expected mean, covariance matrices, and risk aversion coefficient for a mean-variance optimization. For a standard, unconstrained, utility-based optimization, the optimal weights will be:

$$\omega = \frac{1}{\delta}\,\Sigma^{-1}\mu$$

This implies that mean-variance optimization overweights (underweights) those assets that have large (small) estimated returns, negative (positive) correlations, and small (large) variances.[20] Then we can measure the sensitivity of the weights to the expected return by deriving the formula:

$$\frac{\partial\omega}{\partial\mu} = \frac{1}{\delta}\,\Sigma^{-1}$$

This implies that changes in μ tend to lead to large changes in portfolio weights. However, it is very likely to have large estimation errors when estimating the expected returns μ using historical sample data. As a result, those assets with positive pricing errors (estimated return higher than the true expected return) are significantly overweighted versus those with negative errors (estimated return lower than the true expected return).[21] To overcome these problems in mean-variance optimization and incorporate investors' views into this framework, Black and Litterman developed the Black-Litterman model in 1992. The BL model provides not only the equilibrium market portfolio as a starting point for estimation of asset returns, but also a clear way to specify investors' views on returns and to blend the investors' views with prior information. This is made explicit in Figure 2.

Using Bayes' theorem, we can derive the new expected combined return vector as follows:[22]

$$E[R] = [(\tau\Sigma)^{-1} + P^T\Omega P]^{-1}[(\tau\Sigma)^{-1}\Pi + P^T\Omega^{-1}Q]$$

where:

- $E[R]$ is the new combined return vector.
- τ is the weight-on-views scalar.
- Σ is the covariance matrix.
- P is the matrix that identifies the asset involved in the different views.

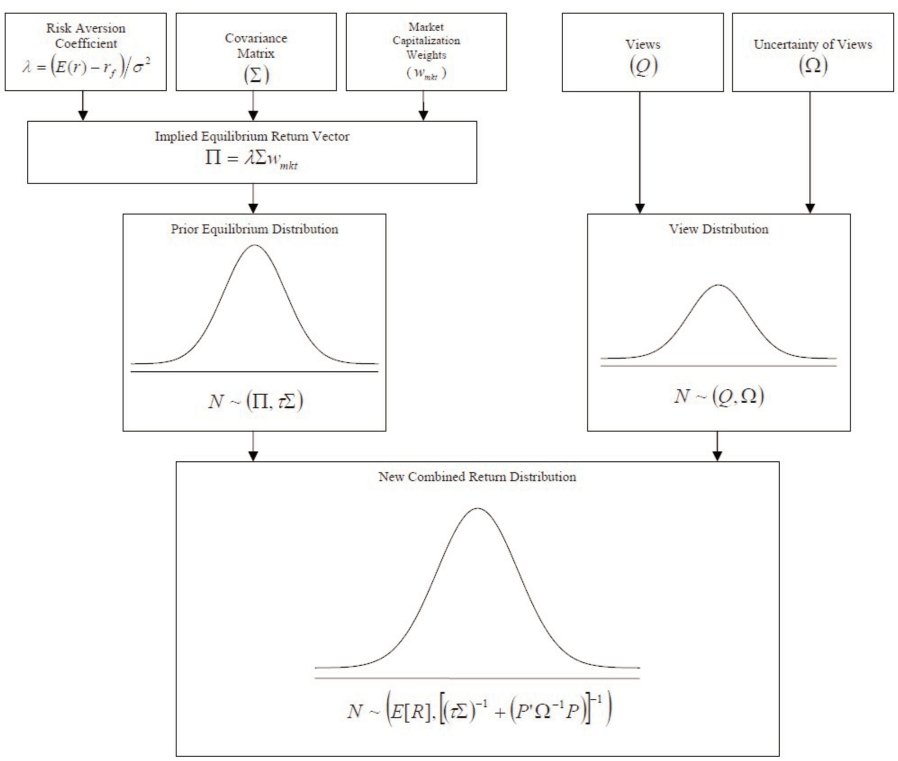

Figure 2 — Black-Litterman model: deriving new expected combined return vector, $E[R]$.

- Ω is a matrix that identifies the uncertainty in the views.
- Π is the implied equilibrium return vector.
- Q is the estimated return vector for every different view.

After we obtain the new combined return vector, we can substitute it into the mean-variance model to obtain the final optimal weights on each asset class.

Downside Risk Minimization (DRM)

Downside risk is the risk of the actual return being below the expected return, or the uncertainty about the magnitude of that difference. Value-at-risk (VaR) measures the potential loss in value of a risky asset or a portfolio over a defined period for a given confidence interval. VaR is defined as the (1 – β)-quantile of the distribution function F of losses of the portfolio; that is:

$$\mathrm{VaR}_{1-\beta} = F^{-1}(1-\beta)$$

where:

$$\beta = \mathrm{P}(\iota > VaR)$$

is called the shortfall probability. In other words, there is only a probability of β that losses can exceed $\mathrm{VaR}_{1-\beta}$.[23]

Betterment mainly uses Black-Litterman to build personalized portfolios for investors.[24] But with different time horizons, investors have different risk tolerance. To balance risk and time, Betterment first predicts the possible return distributions for one portfolio at a given level of risk (i.e., a 70% stocks portfolio) under different market scenarios (from top 5% to bottom 5% of potential outcomes) over 0-20 years.[25] Then Betterment simulates returns for its portfolios at different risk levels (from 0% to 100% stocks) and compares the performances at a similar percentile (market scenario) over a fixed time horizon. Under different percentiles, Betterment can have the best stock allocation curve over different time horizons; thus, Betterment can have the best stock allocation under all percentiles. Finally, Betterment calculates the average "best" stock allocation across all percentiles to arrive at its goal-specific glide path.[26]

Full-Scale Optimization (FSO)

The normal distribution has a skewness (the third standardized moment) of zero and a kurtosis (fourth standardized moment) of three. However, asset returns probability distributions in general are not normal distribution, but feature both skewness and excess kurtosis. Negative (positive) skew indicates that the tail on the left (right) side of the probability density function is longer or fatter than the right (left) side. The excess kurtosis (kurtosis minus 3) greater than zero is platykurtic. Researchers usually explain investors' preferences of these higher moments (i.e., skewness and kurtosis) with a complex utility function. Full-scale optimization is based on the entire empirical return distribution.[27] In the single period portfolio model:

$$\theta^* = \arg\ max_\theta\ U(\theta' R)$$
$$\theta \in \Omega$$

- $R_{n \times T}$ is the expected returns matrix of n assets in T different scenarios.
- $\theta_{n \times 1}$ is the weights vector for each asset in the portfolio.
- U is the utility function.
- Ω is the constraint matrix.

In the mean-variance approach, the utility function is quadratic, resulting in the mean and variance framework. In full-scale optimization, the bilinear function and S-shaped function are the most widely used (see Table 3).

In Table 3, r_p represents portfolio return. In the bilinear utility function, x is the critical return level called the "kink." P is the penalty level for returns lower than the kink. And in the S-shaped function, z is the critical level. A, B, γ_1, and γ_2 determine the curvature.

In Figure 3, the parameters are, x = z = 0, P = 5, A = 1, B = 2, $\gamma_1 = 0.3$, $\gamma_2 = 0.7$. As is shown in the graphs, the zero-return is an inflection point for the utility functions. The two utility functions imply risk aversion, and the downside risk is always penalized. The investor's preference of loss aversion is incorporated in full-scale optimization.

Schwab Intelligent Portfolios uses full-scale optimization in its asset allocation. From the given resources, the bilinear utility function is most likely to be used in Schwab Intelligent Portfolios. And in the optimization, the return threshold is also set to 0; the slope below the threshold is set to two, which means that the pain of losses is two times more painful than the joy an investor would experience from a similar-sized gain.[28]

Utility Function	Utility	
Bilinear	$\ln(1 + r_p)$	for $r_p \geq x$
	$P(r_p - x) + \ln(1 + x)$	for $r_p < x, P > 0$
S-Shaped	$-A(z - r_p)^{\gamma_1}$	for $r_p \leq z$
	$+B(z - r_p)^{\gamma_2}$	for $r_p > z$

Table 3 — Functions used in full-scale optimization.

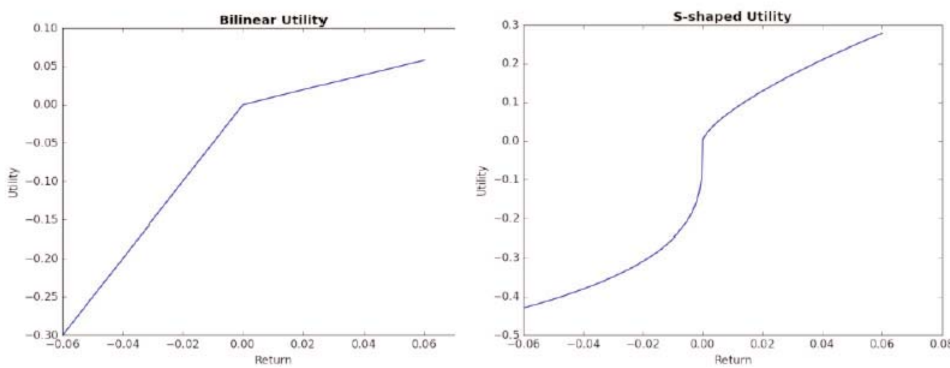

Figure 3 — Utility function.

Monitoring and Rebalancing

Yale University's chief investment officer, David F. Swensen, found that rebalanced portfolios earned an average of 0.4% more per year, with less risk, over 10 years, than portfolios that were not rebalanced.[29] Economist Burton Malkiel reported similar results over a different 10-year period.[30] Robo-advisor companies generally employ threshold-based rebalancing, instead of time-based rebalancing, to take advantage of market movements.

To monitor portfolios, Schwab Intelligent Portfolios uses the Goal Tracker to help investors set up and track their investment goals. The Goal Tracker uses long-term expected return estimates and Monte Carlo simulations to show investors a better, average, and worse outcome, thus monitoring whether their portfolios are "on target" and making necessary adjustments to reach their goals. To create these long-term return estimates, Charles Schwab Investment Advisory (CSIA) uses a set of factors including interest rates, earnings, dividends, and others to estimate future returns and risk on a wide range of asset classes.[31] A Monte Carlo simulation uses repeated sampling to determine the properties of some phenomenon (or behavior).[32] To create the simulations, Goal Tracker uses long-term return estimation to estimate the performance of investors' portfolios.

Schwab Intelligent Portfolios combines rebalancing and tax-loss harvesting as part of a unified process to coordinate the impact of trades. Here we just look at its rebalancing algorithm. On a daily basis, the algorithm checks to determine if asset class weights remain within a drift tolerance. If an asset class has drifted above its target allocation by a sufficient amount, the asset class is sold to bring it back to its target allocation. The cash received from the sale is allocated to the asset classes that are underweight. If an asset class has drifted below its target allocation by a sufficient amount, additional shares of the asset class are bought to bring it up to its target allocation. To make this purchase, overweight asset classes are sold down until there is enough cash to make the purchase.[33] Beyond this, Betterment and Wealthfront also use cash inflows to buy underweighted asset classes.

Overall Comparison

Among the online financial advisory firms that leverage automation and algorithms to help manage client portfolios, Betterment, Wealthfront, and Schwab Intelligent Portfolios stand out for their low account minimums, easy-to-use interfaces, and innovative features. Table 4 gives a summary of the main features of the three robo-advisor companies.

	Schwab Intelligent Portfolios	Betterment	Wealthfront
AUM	$5 billion	$5 billion	$3.5 billion
Management Fee	0%	0.15% to 0.35%	0.25%, $ 0 for first $10,000
Account Minimum	$5,000	$0	$500
Asset Class	5/21	2/14	4/6
ETFs	30	29	21
Goals	6	4	3
Allocation Model	MV & FSO	BL & DRM	MV & CAPM & BL

MV = mean variance; **FSO** = full-scale optimization; **BL** = Black-Litterman; **DRM** = downside risk minimization;
CAPM = capital asset pricing model

Table 4 — Robo-advisors: overall comparison.

Outlook

Compared to traditional financial services, robo-advisors not only provide automated, algorithm-based portfolio management advice, but also charge much less and have much lower minimum requirements for investors. As a result, A.T. Kearney predicts the compound annual growth rate of the robo-advisor market will reach 68% in the next five years.[34]

In China, robo-advisors are still in an emerging phase since their introduction in 2014. Most of the existing products were launched in 2015, and they are much smaller than those abroad. These products can be divided into three categories, according to the companies that they are subject to:

1. Traditional financial company (e.g., Ping An, Harvest Fund)
2. Independent third party (e.g., MiCai, Clipper Advisor)
3. Internet company (e.g., JD Finance, Snow, RoyalFlush)

When allocating asset classes, the mean-variance model, CAPM, BL model, and some other newly devolved models are popular among these robo-advisors in China. This is just like what their foreign peers are doing, but investors never know the details. For example, the way subjective views are generated is really a puzzle when a robo-advisor uses the BL model. This problem, of course, also exists in robo-advisors in the US. At the same time, robo-advisors in China mainly invest in A-shares, public funds, and Internet financial products, which goes against the basic passive index investing methodology of robo-advisors abroad.

Based on the passive index investing methodology, we think the future development of robo-advisors both in China and abroad lies in improving the mechanism for generating and rebalancing subjective views. Radford University finance professor Steven L. Beach and US Securities and Exchange Commission (SEC) economist Alexei G. Orlov have applied EGARCH-M models to generate views[35] and have obtained satisfying results because these financial time series models capture many regularities of stock returns in an elegant and systematic way. Furthermore, machine learning techniques can be applied to forecast expected returns and volatilities. Machine learning models like support vector machines and neural networks have been adopted in volatility prediction by many recent research efforts. These techniques can be incorporated with the BL framework to improve performance. For rebalancing portfolios, an optimal strategy needs to combine threshold-based rebalancing and cash-in, flow-based rebalancing to reduce trading cost and improve performance.

Endnotes

[1]"Robo-advisor." Wikipedia (https://en.wikipedia.org/wiki/Robo-advisor).

[2]Chishti, Susanne, and Janos Barberis. *The FINTECH Book: The Financial Technology Handbook for Investors, Entrepreneurs, and Visionaries*. Wiley, 2016.

[3]"Digital Wealth Management Market Update: A Mosaic of Models Emerges." Aite Group, 3 March 2015 (http://aitegroup.com/report/digital-wealth-management-market-update-mosaic-models-emerges).

[4]Epperson, Teresa, et al. "Hype vs. Reality: The Coming Waves of 'Robo' Adoption." A.T. Kearney, June 2015 (www.atkearney.com/documents/10192/7132014/Hype+vs.+Reality_The+Coming+Waves+of+Robo+Adoption.pdf).

[5]"Financial Engines Reports First Quarter 2016 Financial Results." Financial Engines, 10 May 2016 (https://financialengines.com/education-center/financial-engines-reports-first-quarter-2016-financial-results/).

[6]Epperson et al. (see 4).

[7]Lee, Raef. "5 Ways Robo-Advisors Will Change the Way Advisors Work." SEI Investments Management Corporation, 26 September 2013 (http://seicblogs.com/advisor-practice-management/client-service/5-ways-robo-advisors-will-change-the-way-advisors-work/).

[8]"Schwab Intelligent Portfolios Asset Allocation White Paper." Schwab Wealth Investment Advisory, Inc., 2017 (https://intelligent.schwab.com/public/intelligent/insights/whitepapers/asset-allocation.html).

[9]Bogle, John C. "The Economic Role of the Investment Company" (Senior Thesis). Department of Economics and Social Institutions, Princeton University, 1951.

[10]"Selecting Exchange-Traded Funds for Schwab Intelligent Portfolios." Schwab Wealth Investment Advisory, Inc., 2017 (https://intelligent.schwab.com/public/intelligent/insights/whitepapers/selecting-etf-funds.html).

[11]"Our Investment Selection Methodology." Betterment (www.betterment.com/resources/research/etf-portfolio-selection-methodology/).

[12]"Our Goals and Advice Explained." Betterment (www.betterment.com/resources/research/goals-advice-explained/).

[13]Roszkowski, Michael J., Geoff Davey, and John E. Grable. "Insights from Psychology and Psychometrics on Measuring Risk Tolerance." *Journal of Financial Planning*, Vol. 18, No. 4, April 2005 (www.onefpa.org/journal/Pages/Insights%20from%20Psychology%20and%20Psychometrics%20on%20Measuring%20Risk%20Tolerance.aspx).

[14]"Schwab Intelligent Portfolios Investor Profile Questionnaire White Paper." Schwab Wealth Investment Advisory, Inc., 2017 (https://intelligent.schwab.com/public/intelligent/insights/whitepapers/investor-profile-questionnaire.html).

[15]Barber, Brad M., and Terrance Odean. "Boys Will Be Boys: Gender, Overconfidence, and Common Stock Investment." *Quarterly Journal of Economics*, Vol. 116, No. 1, February 2001 (https://academic.oup.com/qje/article-abstract/116/1/261/1939000/Boys-will-be-Boys-Gender-Overconfidence-and-Common?redirectedFrom=fulltext).

[16]"Lagrange multipliers." Wikipedia (https://en.wikipedia.org/wiki/Lagrange_multiplier).

[17]Merton, Robert C. "An Analytic Derivation of the Efficient Portfolio Frontier." *Journal of Financial and Quantitative Analysis*, Vol. 7, No. 4, September 1972 (www.jstor.org/stable/2329621?seq=1#page_scan_tab_contents).

[18]Black, Fischer, and Robert Litterman. "Global Portfolio Optimization." *Financial Analysts Journal*, Vol. 48, No. 5, September-October 1992 (www.cfapubs.org/doi/abs/10.2469/faj.v48.n5.28).

[19]Fama, Eugene F., and Kenneth R. French. "The Capital Asset Pricing Model: Theory and Evidence." *Journal of Economic Perspectives*, Vol. 18, No. 3, Summer 2004 (www.aeaweb.org/articles?id=10.1257/0895330042162430).

[20]Michaud, Richard O. "The Markowitz Optimization Enigma: Is Optimized Optimal?" *Financial Analysts Journal*, Vol. 45, No. 1, January-February 1989 (www.cfapubs.org/doi/pdf/10.2469/cp.v1989.n4.6).

[21]Christodoulakis, George A. "Bayesian Optimal Portfolio Selection: The Black-Litterman Approach." Notes for Quantitative Asset Pricing, MSc Mathematical Trading and Finance, November 2002 (www.globalriskguard.com/resources/assetman/bayes_0008.pdf).

[22]Christodoulakis (see 21).

[23]Gilli, Manfred, Evis Këllezi, and Hilda Hysi. "A Data-Driven Optimization Heuristic for Downside Risk Minimization." *Journal of Risk*, Vol. 8, No. 3, Spring 2006 (https://papers.ssrn.com/sol3/papers.cfm?abstract_id=900365).

[24]"Portfolio Optimization: Our Secret to Driving Better Performance." Betterment (www.betterment.com/resources/investment-strategy/portfolio-management/portfolio-optimization/).

[25]"Our Stock Allocation Advice." Betterment (www.betterment.com/resources/research/stock-allocation-advice/).

[26]Betterment (see 25).

[27]Hagströmer, Björn, et al. "Mean-Variance vs. Full-Scale Optimization: Broad Evidence for the UK." Federal Reserve Bank of St. Louis, Working Paper No. 2007-016D, April 2007 (https://research.stlouisfed.org/wp/2007/2007-016.pdf).

[28]Schwab (see 8).

[29]Swensen, David F. *Unconventional Success: A Fundamental Approach to Personal Investment*. Free Press, Simon & Schuster, 2005.

[30]Malkiel, Burton G., and Charles D. Ellis. *The Elements of Investing: Easy Lessons for Every Investor*. John Wiley & Sons, 2012.

[31]"Goal Tracker for Schwab Intelligent Portfolios." Schwab Wealth Investment Advisory, Inc., 2017 (https://intelligent.schwab.com/public/intelligent/insights/whitepapers/goal.html).

[32]Sawilowsky, Shlomo S., and Gail C. Fahoome. "Statistics via Monte Carlo Simulation with Fortran." *Journal of Modern Applied Statistical Methods*, 2003 (https://en.wikipedia.org/wiki/Shlomo_Sawilowsky).

[33]"Rebalancing and Tax-Loss Harvesting in Schwab Intelligent Portfolios." Schwab Wealth Investment Advisory, Inc., 2017 (https://intelligent.schwab.com/public/intelligent/insights/whitepapers/tax-loss-harvesting-rebalancing.html).

[34]Epperson et al. (see 4).

[35]Beach, Steven L., and Alexei G. Orlov. "An Application of the Black-Litterman Model with EGARCH-M-Derived Views for International Portfolio Management." *Financial Markets and Portfolio Management*, Vol. 21, No. 2, June 2007 (https://link.springer.com/article/10.1007/s11408-007-0046-6).

Jie Yang, Hanxi Ye, Yadan Wei, and Linqian Bao are Business System Analysts for the Wangfubao Project in State Street Technology (Zhejiang). Since 2010, Wangfubao has been dedicated to helping retail investors in China plan, save, and invest in the mutual fund market to reach their financial goals. Wangfubao tries to combine interactive visual analysis methods for high-dimensional time series data, collective intelligence of the investor community, and modern portfolio theory to build investment strategies and do asset allocation. The authors can be reached at JYang7@StateStreet.com, HYe@StateStreet.com, YWei@StateStreet.com, and LBao@StateStreet.com.

Appendix: Investor Profile Questionnaires

Schwab Intelligent Portfolios Questionnaire

1. My goal for this account is to:
   a. Prepare for retirement
   b. Save for major upcoming expenses (education, healthcare bills, etc.)
   c. Save for something special (vacation, new car, etc.)
   d. Build a rainy day fund for emergencies
   e. Generate income for expenses
   f. Build long-term wealth
2. I have ___ understanding of stocks, bonds, and ETFs.
   a. No
   b. Some
   c. Good
   d. Extensive
3. When I hear "risk" related to my finances:
   a. I worry I could be left with nothing.
   b. I understand that it's an inherent part of the investing process.
   c. I see opportunity for great returns.
   d. I think of the thrill of investing.
4. Have you ever lost 20% or more of your investments in one year?
   a. Yes
   b. No
5. In the year I lost 20% of my investments/ If I ever were to lose 20% or more of my investments in one year, I would:
   a. Sell everything
   b. Sell some
   c. Do nothing
   d. Reallocate my investments
   e. Buy more
6. When it comes to making important financial decisions:
   a. I try to avoid making decisions.
   b. I reluctantly make decisions.
   c. I confidently make decisions and don't look back.
7. I am ___ years old.
8. My initial investment for this goal is ___.
9. One year from now, I'd be comfortable with my initial investment fluctuating between: (indicate range around initial investment size).
10. I plan to save an additional ___ per month for this goal.
11. I need the money for this goal starting in x years for y years. (Specify x and y.) OR I need income for x years ("Generate income for expenses" goal).

Betterment Questionnaire

1. I am ___ years old.
2. I am retired/not retired.
3. My annual income is ________.
4. What's my plan? (retirement, safety net, general investing, major purchase)

Wealthfront Questionnaire

1. What's your primary reason for investing?
   a. General savings
   b. Retirement
   c. Other

2. What are you looking for in a financial advisor? (Select all that apply.)
    a. I'd like to create a diversified investment portfolio.
    b. I'd like to save money on my taxes.
    c. I'd like someone to completely manage my investments, so that I don't have to.
    d. I'd like to match or beat the performance of the markets.
3. What is your current age?
4. What is your annual pre-tax income?
5. Which of the following best describes your household?
    a. Single income, no dependents
    b. Single income, at least one dependent
    c. Dual income, no dependents
    d. Dual income, at least one dependent
    e. Retired or financially independent
6. What is the total value of your cash and liquid investments (e.g., savings, CDs, mutual funds, IRAs, 401[k]s, public stocks)?
7. When deciding how to invest your money, which do you care about more?
    a. Maximizing gains
    b. Minimizing losses
    c. Both equally
8. The global stock market is often volatile. If your entire investment portfolio lost 10% of its value in a month during a market decline, what would you do?
    a. Sell all of your investments
    b. Sell some
    c. Keep all
    d. Buy more

Going Global: Internationalizing and Localizing a Legacy Financial System

by Bo Zhou and Lucy Chen

Overview

Nowadays, software applications need to be distributed to different countries and regions across the world. It is common for a popular software package to have many local versions; for example, Windows 7 has a US version, French version, simplified Chinese version, and so on. The local versions of a software application help local users better understand and use it, attract more users, and increase software sales.

Internationalization and localization are important steps in distributing and deploying systems to different regions of the world. Internationalization refers to the process of engineering a system such that it can support various languages and regions without further modification. Localization refers to the process of adapting an internationalized software system for a specific language or region. Unfortunately, not all software products were originally designed with multi-language support. Some legacy systems developed between the 1980s and 1990s did not take software internationalization and localization into consideration, yet for various reasons they are still in service around the world.

A couple of books have been published on software internationalization[1] and localization.[2] These books are good sources of information on these subjects and introduce many technologies and tools. However, in a real commercial project, the process of software internationalization and localization is a bit different in terms of how to use such technologies and tools, and these books report no industrial experience. In this article, we present our experience in, and propose a process and tool supports for, software internalization and localization.

The system we reengineered, HOLLANDER, is a large-scale legacy financial system at State Street Corporation that contains more than 5 million lines of C/C++ source code and 30 different modules grouped into clusters. The whole project lasted for two years, from 2009 to 2011, and involved 25 developers and eight quality assurance (QA) personnel. We iterated the reengineering work from cluster to cluster. For each iteration, we used our source code ranker, "IRanker," to find the most important source files. We extracted some patterns from these files and used them in our automated code search tool "I18nLocator." These patterns were divided into two categories: convertible patterns and suspicious patterns. The I18nLocator automatically converted the source code located by convertible patterns and highlighted the source code located by suspicious patterns.

Taking into account the context of the lines of source code (SLOC) located by suspicious patterns, developers reengineered them. There might be some patterns that we missed by just analyzing the important files highlighted by IRanker. So we ran unit tests and performed other internal QA activities to find bugs caused by the missing patterns. Additional patterns to address these bugs were then identified. After we internationalized HOLLANDER, we extracted hard-coded strings and put these in a separate resource data file.

HOLLANDER: A Brief Introduction

HOLLANDER was developed in the 1990s, and it has a traditional client-server architecture (see Figure 1). HOLLANDER uses Microsoft Foundation Class (MFC) technology to build its client and socket communication to transfer messages between client and server. The architecture contains three parts:

1. The **application server** handles messages from the client and decides whether to communicate with data servers to generate the required data.
2. **Data servers** mainly fetch data from a database or execute some operations on the database (such as add, modify, remove, and delete records) according to the requests from the application server.

3. The **database** stores and persists data. The four layers of our client-server architecture could be mapped to the MFC framework layer, application layer, data access layer, and database layer, as shown in Figure 1.

Figure 2 presents the 30 different modules in HOLLANDER. According to the functionalities of these modules, we group them into six clusters.

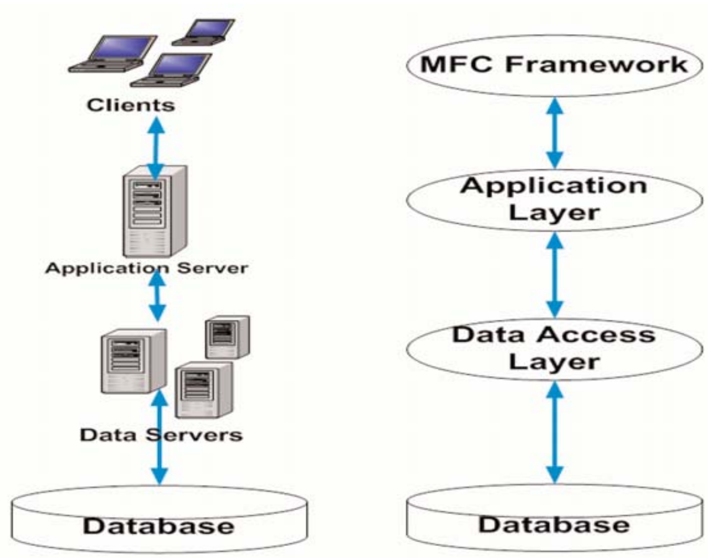

Figure 1 — The architecture of HOLLANDER.

Overall Framework

Software internationalization and localization refer to two different processes in reengineering HOLLANDER. Localization depends on the internationalized version of HOLLANDER. Figure 3 presents our proposed internationalization and localization framework. We believe that this framework could be used for other internationalization and localization projects.

We began by extracting important source code files using our source code ranker IRanker (Step 1). Then we analyzed the important SLOC to extract convertible patterns and suspicious patterns (Steps 2 and 3). Next we input these patterns into I18nLocator (Step 4), an automated tool for locating the SLOC that match these patterns. I18nLocator automatically converts the SLOC that match convertible patterns and highlights the lines that match suspicious patterns so developers can perform manual conversion (Step 5). After the above reengineering steps, we also performed unit testing and local QA checks (Step 6). We analyzed the failures captured by the unit tests and QA checks, discovered the patterns we missed, and removed/edited the incorrect patterns (Step 7). We again located these patterns using I18nLocator (Step 8). By iterating the above steps many times for each cluster of HOLLANDER modules, we finally succeeded in internationalizing HOLLANDER.

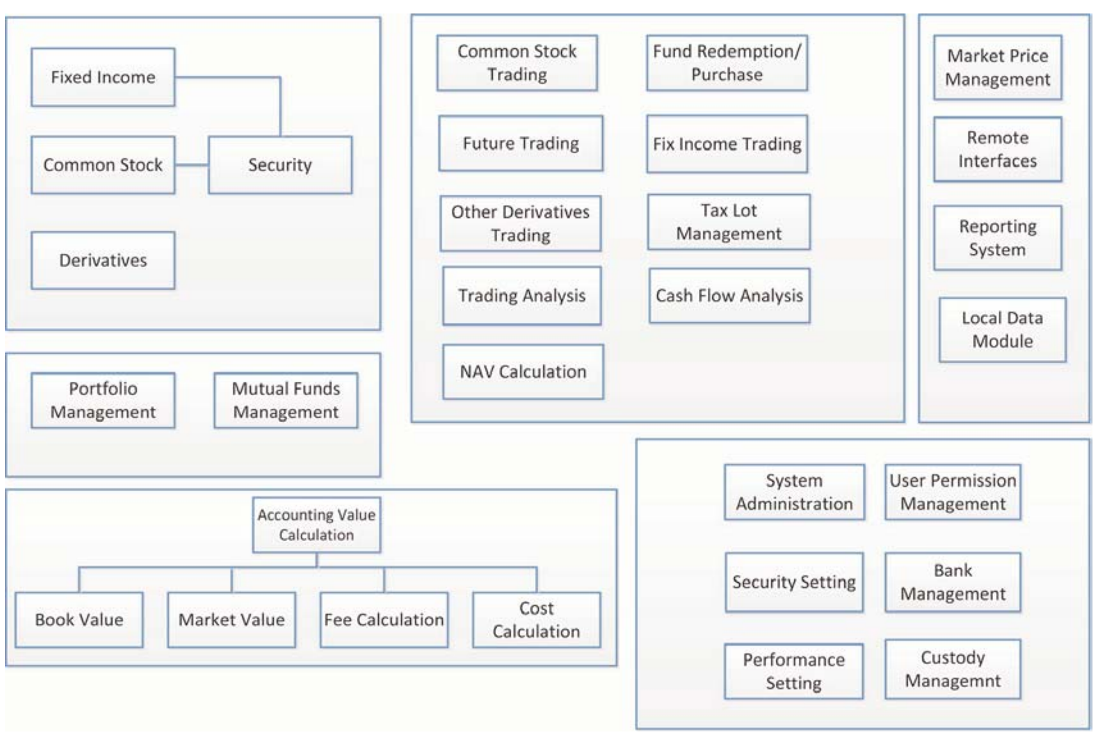

Figure 2 — HOLLANDER modules.

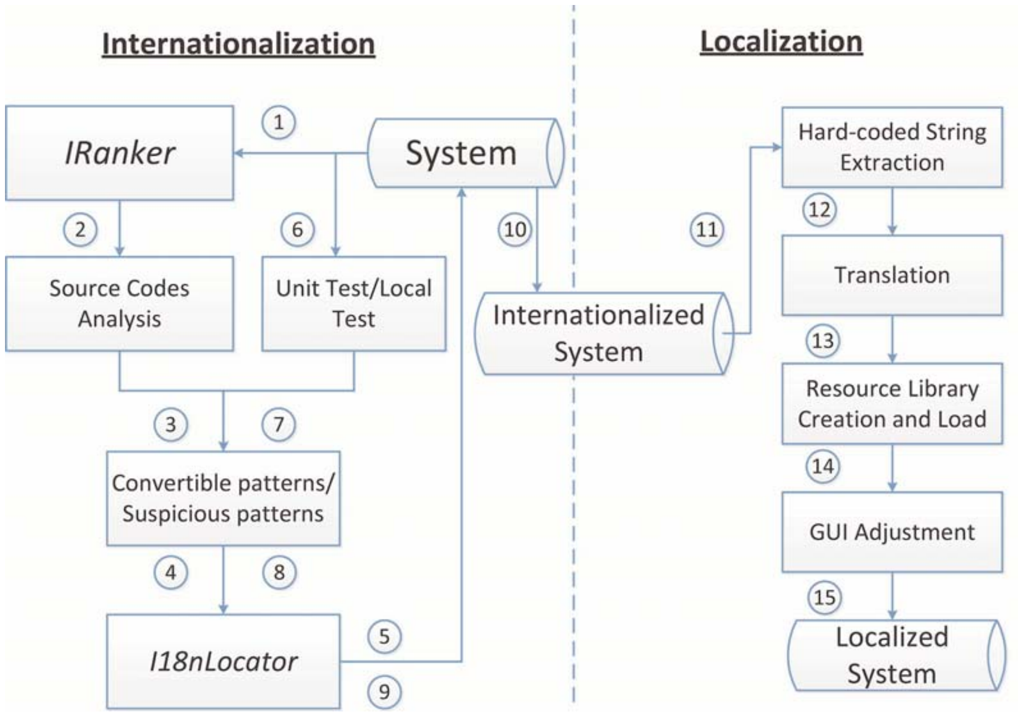

Figure 3 — The overall internationalization and localization framework.

For HOLLANDER localization, we first extracted the hard-coded strings from the SLOC (Step 11). Then we translated these hard-coded strings to Chinese (Step 12), and for each of the different languages (i.e., English and Chinese), we provided different resource library files to cleanly separate resource data from source code (Step 13). Finally, we adjusted the GUI widgets to ensure that they have enough space to display text in different languages (Step 14).

Software Internationalization Methodology

General Idea

The main task for system internationalization is Unicode transformation; that is, to replace the types and functions in the SLOC that do not support the Unicode standard to Unicode-compatible ones. Concretely, Unicode transformation includes the transformation of the following five program element types:

1. Character variable
2. Function
3. Windows macro
4. Database field
5. Text processor

IRanker: Ranking Important Code

Large-scale legacy systems, such as HOLLANDER, do not have many documents that describe them, and as we discuss below, we need to find the transformation patterns from a small yet representative set of SLOC. In this section, we discuss IRanker, which can be used to detect a representative set of code for manual identification of convertible and suspicious patterns.

First we build a dependency graph[3] at the source code file level; we use a directed graph to represent the dependency relationships among source files. For example, for any two source files A and B, if A uses any of the methods in B, then A and B have a dependency relationship, and there is an edge from A to B. In the C/C++ language, the dependency is defined at the top of source files, by statements like "include <stdio.h>". We use these "include" statements to build a dependency graph among source files, and we use the corresponding source code files to replace the header files referred to in the "include" statements. For example, in source file "A.c", it includes "B.h", and in our dependency graph, we just use "B.c" to replace "B.h". We remove the standard C/C++ libraries and other libraries (e.g., Oracle libraries, MFC libraries) as we do not want to reengineer them.

Convertible Patterns vs. Suspicious Patterns

We represent these patterns by regular expressions. We refer to the patterns that identify the SLOC that can be automatically converted to their corresponding Unicode-compatible counterparts as convertible patterns. We refer to the patterns that identify the SLOC that require developers' judgment as suspicious patterns.

Figure 4 summarizes the process of extracting convertible and suspicious patterns. Suspicious patterns require developer inspection, as the modifications of the corresponding SLOC need to consider the context of the SLOC. Often the root cause of the modifications corresponding to the suspicious patterns is that the size of a character increases from 1 byte to 2 bytes, which causes the memory space requirement of related data structures and operations to change.

We consider five suspicious pattern types:

1. **String pointer.** In C/C++, we use a pointer to store the memory address of a variable. The SLOC that compute the address using the pointer and its offset address are suspicious code.
2. **Function.** There are three types of suspicious patterns related to function definitions and invocations: (1) no corresponding Unicode-compatible function, (2) parameter difference between corresponding Unicode and ASCII functions, and (3) parameter misuse.
3. **COM component.** Some COM-related API functions only accept wide characters, and in the original HOLLANDER, developers used "MultiByteToWideChar" and "WideCharToMultiByte" functions to do the transformations between wide characters and single-byte characters.
4. **File operation.** To internationalize HOLLANDER, we needed to make all file operation-related codes support Unicode (mainly UTF-8, UTF-16) and ASCII. We developed a new file operation class UniFileOperation that supports Unicode and ASCII, located the SLOC related to file operations, and replaced these SLOC with the invocations of appropriate functions in the UniFileOperation class. The internal process of UniFileOperation is presented in Figure 5.
5. **Third-party component.** In HOLLANDER, there are many third-party components. For these components, we do not have any source code. One way to internationalize them is to update them to the latest versions.

I18nLocator

I18nLocator is an automated tool we use to locate the SLOC that match suspicious patterns and convert the SLOC that match convertible patterns. It contains two components, the replacer and detector:

1. A **replacer** component mainly converts the SLOC that match convertible patterns.
2. A **detector** component is used to locate suspicious SLOC that match suspicious patterns. These suspicious patterns are also represented by regular expressions. Developers use I18nLocator to locate suspicious SLOC in HOLLANDER. Whenever suspicious code is detected, I18nLocator will add a

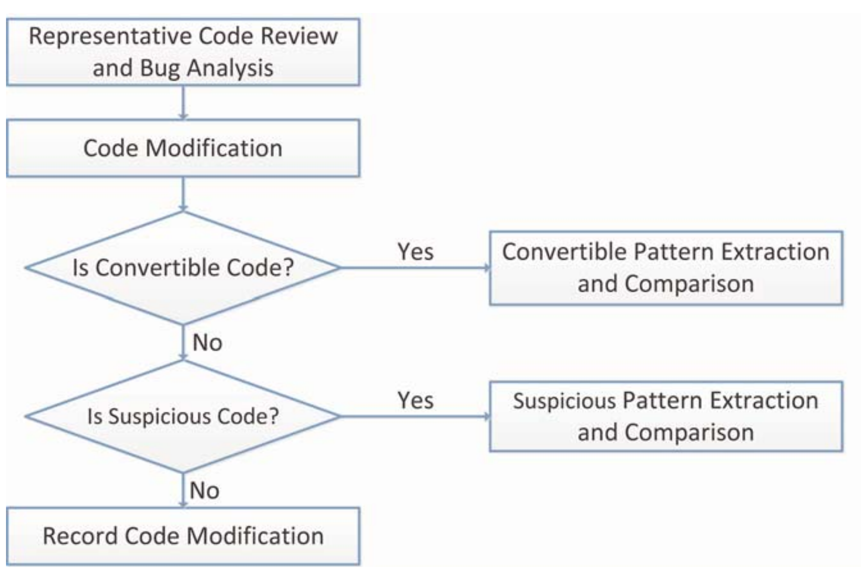

Figure 4 — Suspicious and convertible patterns extraction process.

comment behind these SLOC. The format of the comment is: //SUSPICIOUS(OPEN): {Corresponding Regular Expressions}. When I18nLocator is run again, it will ignore the SLOC with comments "SUSPICIOUS(CLOSED)".

Figure 6 presents the GUI of the detector component in I18nLocator.

Software Localization Methodology

Software localization requires us to adapt internationalized software for a specific region or language. In this section, we describe our steps for HOLLANDER localization.

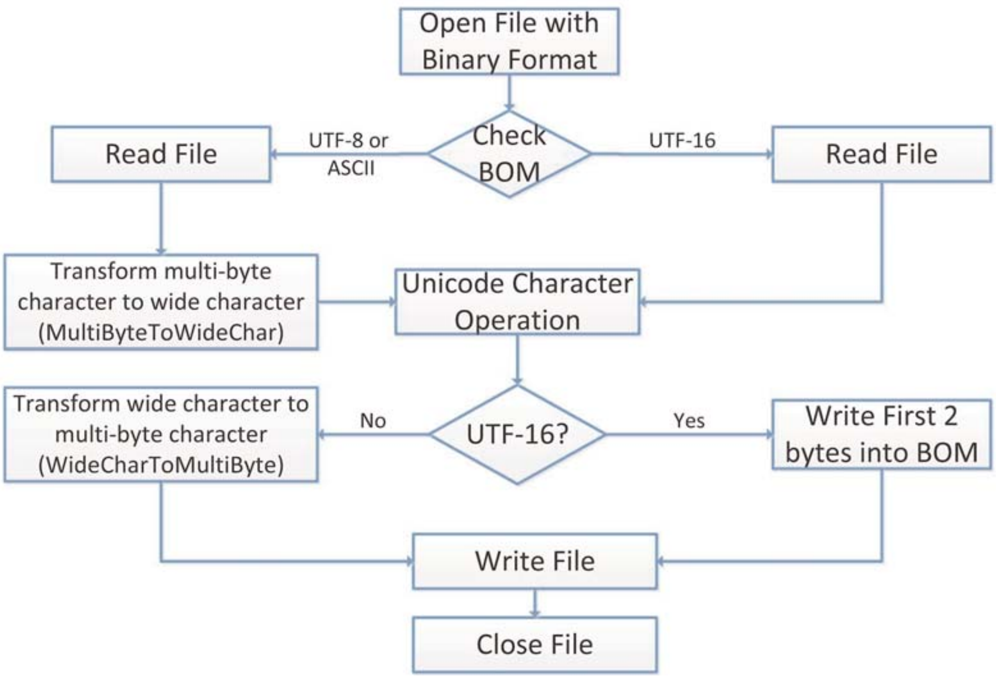

Figure 5 — File operation process of UniFileOperation class.

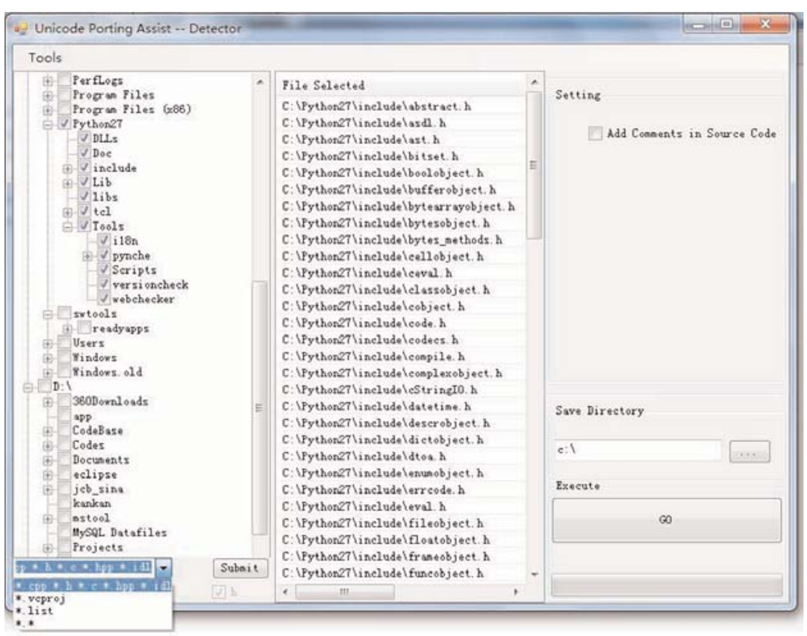

Figure 6 — GUI of detector component of I18nLocator.

Hard-Coded String Extraction

Hard-coded strings refer to strings that can be displayed in the GUI, and they are directly written in the SLOC. An example is:

```
1: CString loginFailedMsg = "the password is wrong";
```

For a hard-coded string, we need to extract it, use a variable to replace it, and store the mapping relationship between the variable and hard-coded string in a resource data file. Unlike SLOC targeted by our internationalization methodology, hard-coded strings have a single pattern; that is, they begin with a " or a ', end with a " or a ', and do not appear in macro definition statements. We use a regular expression to represent this pattern. A typical process of hard-code string extraction and storage in a resource data file is presented in Figure 7. Different languages have different resource data files. To localize HOLLANDER, we configured the language setting in a configuration file.

Translation

Translation includes GUI text translation and online help document translation. We outsource these need-to-translate materials (i.e., hard-coded strings and documents) to a third-party local language translation provider. Four types of objects need to be translated:

1. The resource data file, which is used to store hard-coded strings
2. The online help document
3. Constant strings in the database
4. Other text files in HOLLANDER

Translation may cause various bugs in HOLLANDER, and these bugs are hard to fix; we cannot either prevent or ameliorate these bugs.

Our Results

In this section, we report our assessment on the effectiveness of our proposed methodologies for HOLLANDER.

Effectiveness of I18nLocator

Since both IRanker and unit tests/local checks discover convertible and suspicious patterns, we divide the effectiveness of I18nLocator into two parts:

1. IRanker
2. Unit tests/local checks

Table 1 presents the results of an I18nLocator code search using IRanker and unit tests/local checks. Using IRanker, we detected most of the convertible patterns and suspicious patterns; that is, 30 convertible patterns and 25 suspicious patterns, respectively, which cover 78.95% of all the convertible patterns and 71.43% of all the suspicious patterns.

Due to the security confidentiality policy of State Street Corporation, we can only list the number of SLOC to the nearest 1,000. Using the patterns identified by analyzing the source code selected by IRanker, I18nLocator

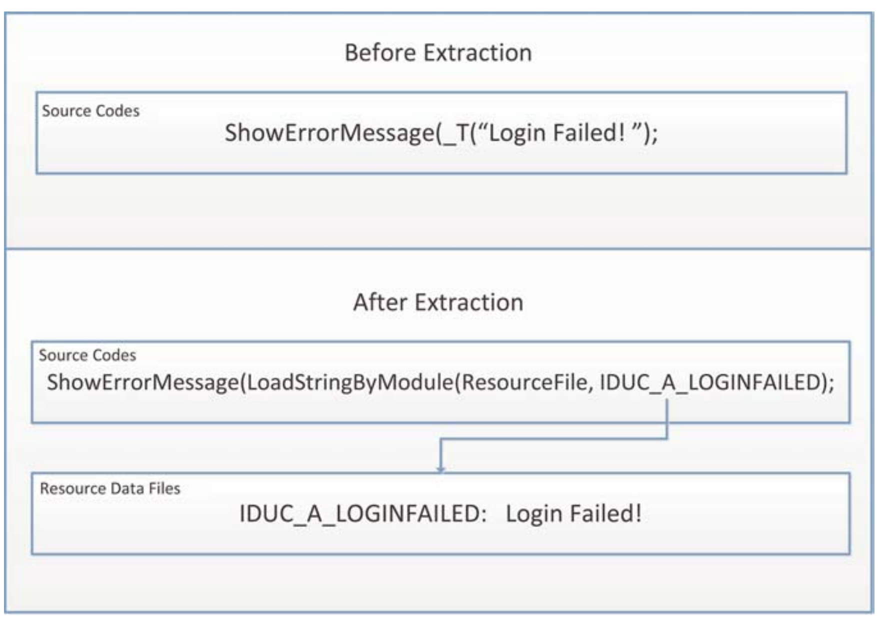

Figure 7 — Hard-coded string extraction and resource data creation process.

Type	Convertible Patterns	Suspicious Patterns	Total Patterns	Convertible LOC	Suspicious LOC	Total LOC
IRanker	30	25	55	350,000	80,000	430,000
Unit tests/ local checks	8	10	18	52,000	22,000	74,000
Total	38	35	73	402,000	102,000	504,000

Table 1 — Results of the I18nLocator code search (all LOC rounded down to the nearest thousand).

automatically converted around 350,000 SLOC that matched convertible patterns and located around 80,000 SLOC that matched suspicious patterns. This covers 87.06% and 78.43%, respectively, of the total source code modification work. IRanker helps HOLLANDER internationalization significantly since it extracts a relatively small proportion of the code from which general patterns can be inferred. These patterns cover most of the changed code. This reduces the testing and analysis work needed.

With the help of IRanker and unit tests/local checks, we discovered 73 patterns and modified 504,000 SLOC. Among those 504,000 lines, we automatically converted 402,000 SLOC, which covers 79.76% of the reengineering work. The remaining SLOC were flagged by the suspicious patterns. Although IRanker and unit tests/local checks do not automatically reengineer the code, they significantly aid developers in locating these SLOC, which is not an easy job due to the complexity of HOLLANDER.

Internationalization and Localization Test

The test for the quality of system internationalization and localization should consider two aspects. First, we should not affect the original functionality of the system. Second, we should support multi-language input and have a local language GUI. The test work should cover these two aspects, and we divide test work into three types:

1. Functional test
2. GUI test
3. Translation test

Due to the software management strategy at State Street, we have three environments:

1. Development (DEV)
2. User acceptance test (UAT)
3. Production (PROD)

DEV is used by local QAs and developers; we perform local checks and unit tests in DEV. In this case, we designed 3,014 test cases to test HOLLANDER. After we passed all these test cases, we delivered HOLLANDER to UAT and further to PROD. Only the bugs discovered in UAT and PROD are recorded in the bug tracking system; there were 110 bugs reported. We modified 504,000 SLOC, and the number of bugs found during UAT and PROD was 0.000218 bugs/SLOC. This shows that the quality of the resultant internationalized and localized system is good.

The bug distribution is presented in Table 2. The number of bugs detected by functional tests, GUI tests, and translation tests are 20, 23, and 61, respectively, which corresponds to 18.2%, 20.9%, and 55.5% of the total number of bugs. It is interesting to note that translation bugs were the majority of all 110 bugs detected when HOLLANDER was deployed in UAT and PROD. Moreover, many functional bugs and GUI bugs were also related to translation.

From Table 2, we also notice that there were relatively few internationalization-related bugs. File processing and non-Unicode program transformation are related to internationalization, but there were only 12 bugs in these categories, which is 10.9% of the total number of bugs. This shows that our proposed software internationalization and localization process is effective.

Related and Future Work

To the best of our knowledge, there is limited research work on software internationalization and localization. Software researcher Xiaoyin Wang and his coauthors propose a method that automatically locates need-to-translate constant strings for software internationalization.[4] They first collect APIs related to the GUI, and based on string-taint analysis, they search for need-to-translate constant strings. Finally, they evaluate the performance of their approach using four open source applications: RText, RISK, Art of Illusion, and MegaMek.

Type	Test Aspects	Number of Bugs
Functional test	General functionality	4
	Resource data files	4
	File processing	4
	Non-Unicode program transformation	8
GUI test	Widget layout	8
	Text display	10
	Shortcut key usage	4
	TAB key sequence	1
Translation test	Online help document	10
	GUI translation	25
	Translation correctness	26
Miscellaneous	Other bugs	6

Table 2 — Type, test, and number of bugs for HOLLANDER.

The authors further extend their work by locating these need-to-translate constant strings for Web application internationalization.[5] They propose a flag propagation-based approach that distinguishes strings visible on the browser side from nonvisible strings and evaluate the performance of their approach using three PHP-based Web applications.

Our work is different from the work of Wang et al. in the following respects:

- We propose a holistic, end-to-end process for software internationalization and localization, along with tool supports. Wang et al. solve one problem in the software internationalization process, which is the locating of constant strings. We use a more lightweight solution to locate constant strings.
- We evaluated our proposed process on a commercial system called HOLLANDER, which contains 5 million SLOC. The largest system evaluated by Wang et al. contains only 110,000 SLOC. Also, the system that we analyzed and reengineered is currently deployed and used in the industry.

In the future, we plan to do more internationalization and localization work on other large-scale commercial systems and to learn more from these experiences.

Endnotes

[1]Uren, Emmanuel, Robert Howard, and Tiziana Perinotti. *Software Internationalization and Localization: An Introduction*. Van Nostrand Reinhold, 1993.

[2]Esselink, Burt. *A Practical Guide to Localization*. John Benjamins Publishing, 2000.

[3]Balmas, Françoise. “Using Dependence Graphs as a Support to Document Programs.” *Proceedings of the Second IEEE International Workshop on Source Code Analysis and Manipulation*. IEEE Computer Society, 2002 (http://ieeexplore.ieee.org/document/1134114/); Balmas, Françoise. “Displaying Dependence Graphs: A Hierarchical Approach.” *Journal of Software Maintenance and Evolution: Research and Practice*, Vol. 16, No. 3, May/June 2004 (http://onlinelibrary.wiley.com/doi/10.1002/smr.291/full).

[4]Wang, Xiaoyin, et al. “Locating Need-to-Translate Constant Strings for Software Internationalization.” *Proceedings of the IEEE 31st International Conference on Software Engineering*. IEEE, 2009 (http://ieeexplore.ieee.org/document/5070535/); Wang, Xiaoyin, et al. “TranStrl: An Automatic Need-to-Translate String Locator for Software Internationalization.” *Proceedings of the IEEE 31st International Conference on Software Engineering*. IEEE, 2009 (http://ieeexplore.ieee.org/document/5070554/).

[5]Wang, Xiaoyin, et al. “Locating Need-to-Translate Constant Strings in Web Applications.” *Proceedings of the 18th ACM SIGSOFT International Symposium on Foundations of Software Engineering*. ACM, 2010 (http://dl.acm.org/citation.cfm?id=1882306).

Bo Zhou is a Professor in the College of Computer Science and Technology in Zhejiang University, China. Dr. Zhou's research interests include software engineering, very large information systems, database management systems, and data analytics. He has a strong connection to the industry, and he is the Senior Technology Advisor to State Street Technology (Zhejiang). Dr. Zhou received his PhD in computer science and technology from Zhejiang University in 1996. He can be reached at bzhou@zju.edu.cn.

Lucy Chen is the Managing Director of State Street Technology (Zhejiang), the Department Head of the Japan & China Fund Accounting Platform, State Street Global Exchange Hangzhou IT, and Data^{GX} Hangzhou Team. Her team covers IT business analysis, development, QA testing, implementation, client service, and support, with 150+ members serving clients in China, Japan, North America, and beyond. Ms. Chen has bachelor's and master's degrees in computer science from Zhejiang University and has held PMP certification since 2007. One of her recent work responsibilities is to serve Japanese business from a fund accounting system perspective. In the past three years, her team has successfully migrated two clients from a Japan-dominant accounting platform to their internal application, and a third one is in progress. Meanwhile, her team has also provided fund accounting software service to several Chinese fund management companies. She can be reached at lu.chen@statestreet.com.

Picture This: Using 3D Visual Analytics to Explore Complex Temporal Data

by Jerry Cristoforo, Qiao Huang, Zhiyu Peng, and Xiaohu Yang

Introduction

Nowadays, large-scale data is growing at explosive speed everywhere. While the capacity to collect and store new data is rapidly growing, the ability to analyze these data volumes is increasing at much lower rates.[1] Although it is not easy to analyze such large-scale complex data, it holds the key to gaining insight into patterns, trends, and correlations, which means greater potential value. In the past, researchers have proposed many automatic analysis methods to process and simplify this data, such as principal component analysis, self-organizing maps, and clustering algorithms. Another typical way to handle large-scale data is online analytical processing (OLAP), which partially pre-aggregates data into "cubes" and stores it in a data warehouse. But we still lack an effective and flexible way to present the data or the analysis results intuitively.

In recent years, visual analytics (VA) has been introduced to represent massive multidimensional temporal data with various visual encodings. VA combines automated analysis techniques with interactive visualizations to enable effective understanding, reasoning, and decision making on the basis of very large and complex data sets.[2] A survey delivered by University of Konstanz computer science professor Daniel Keim and his coauthors[3] gives a good overview of the scope, process, and challenges of this field. A more detailed review of all aspects of visual analytics can be found in Keim et al.'s book *Mastering the Information Age — Solving Problems with Visual Analytics*,[4] which is the outcome of a two-year research project. Keim and his colleagues suggest that visual analytics can be applied to many domains, such as finance, health, geography, physics, and security, among others.

There are many well-designed VA systems and tools that are effective in supporting domain analysts' decision making and discoveries of insights through advanced visualization methods and user-centric interactive operations. One of the most popular and widely used commercial tools for visualizing business data is Microsoft Excel, which provides the standard visualization methods for spreadsheet data (e.g., bar charts, columns, line charts, pie charts). However, these visualization methods are challenged when the underlying data model consists of complex ideas that need to be communicated with clarity, precision, and efficiency.

The core premise of our research is that we can extract many times more insight out of data by visualizing it than we can through more traditional means.

Excellence in visual analytics consists in giving the viewer the greatest number of ideas in the shortest time in the smallest space.[5] VA excellence is often found in simplicity of design and complexity of data. The core premise of our research is that we can extract many times more insight out of data by visualizing it than we can through more traditional means. In this article, we present Apsara, a VA system that supports interactive analysis of multidimensional temporal data. We introduce the idea of a novel 3D visualization method that provides a compact universal overview of large-scale data and allows drilling down for further detailed information. The system extends the ability of existing mapping techniques by visualizing domain data based on a 3D geometry enhanced by color, motion, and sound. The interaction design is based on Keim's VA mantra: "Analyze first — Show the important — Zoom, filter, and analyze further — Details on demand."[6] Apsara is designed with customizable visualization that can be adapted to different data models and applied to multiple domains. To evaluate the system, we applied it to a specific domain (financial services) with real-world data, and the results show that the visualization and analysis of multidimensional temporal data provides significant benefits.

Apsara is flexible enough to apply to different domains and to support:

- A novel 3D visualization method in a universal overview
- Large-scale data management
- Multidimensional data visualization
- Temporal data visualization
- Interactive analysis

Before introducing the technical details of Apsara, we first discuss previous work on visual analytics, including visualization methods and VA systems and tools.

Related Work

Visualization Methods

Traditional visualization methods (e.g., 2D line charts, bar charts, tables) are simple and intuitive, and thus are widely used in our daily life. But they are limited in terms of showing large-scale complicated data (see Figure 1). In the last few decades, many sophisticated visualization methods have been proposed, such as scatter plots, heat maps, parallel coordinates, treemaps, Internet maps, spiral graphs, density-based distribution maps, sphere-based maps, theme rivers, and time wheels. These methods are useful when visualizing certain types of data; for example:

- Heat maps are well suited to depicting financial information.
- Spiral graphs are a successful method for discovering the periodic pattern of disease outbreaks.
- Parallel coordinates are a common way of visualizing high-dimensional geometry and analyzing multivariate data.

The basic theory of information visualization is summarized in Edward Tufte's book *The Visual Display of Quantitative Information*.[7]

Systems and Tools

A variety of companies — ranging from specialized data discovery vendors such as Tableau, QlikTech, and Spotfire to multinational corporations such as IBM, Microsoft, Oracle, and SAP — have devoted much effort to developing their own commercial VA system products for analyzing data of increasing volume and variety that arrives ever more quickly.[8] However, some big software vendors tend to focus only on a small number of "standard" visualization techniques, such as line charts and tables, which have limited capability in handling large, complex data sets. Open source alternatives include such toolkits as Prefuse, Protovis, and InfoVis Toolkit.[9]

On the academic side, a number of VA systems have also been developed to support domain analysts' work. MobiVis is a system for visually analyzing mobile data

Year	Equity	Fixed-Income	Close-End	QDII	Total
2002	0.04	0.00	0.07	-	0.12
2003	0.07	0.02	0.09	-	0.17
2004	0.16	0.08	0.08	-	0.33
2005	0.17	0.22	0.08	-	0.47
2006	0.59	0.11	0.16	-	0.86
2007	2.74	0.19	0.24	0.11	3.28
2008	1.23	0.59	0.07	0.05	1.94
2009	2.11	0.36	0.12	0.07	2.67
2010	1.99	0.30	0.13	0.07	2.50
2011	1.52	0.46	0.12	0.06	2.17
2012	1.58	1.02	0.13	0.06	2.80

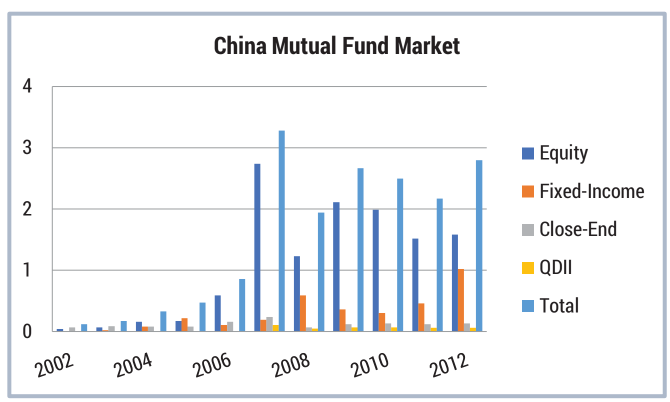

Figure 1 — Visualization of 10 years of data on the Chinese mutual fund market. It is difficult for analysts to find patterns and trends in the table, which only contains a bunch of numbers. The 2D bars in the chart are more suitable for showing past and future trends, but this representation does not allow users to access detailed information on certain types of mutual funds for further exploration.

with the idea of presenting social and spatial information in one heterogeneous network.[10] VIS-STAMP is a visual inquiry system for space-time and multivariate patterns. It supports the overview of complex patterns and, through a variety of interactions, enables users to focus on specific patterns and examine detailed views.[11] WireVis is a system that combines multiple visualization methods to analyze categorical, time-varying data from financial transactions.[12] Weijia Xu, group lead for the Data Mining & Statistics Group at Texas Advanced Computing Center, designed a system based on treemaps to analyze large digital collections with interactive visualization.[13] And Microsoft researcher Danyel Fisher has developed Hotmaps, based on heat maps, to visualize geography data. It takes advantage of the structure of the underlying data set to visualize it in its own space.[14]

Different from their work, our system is more flexible and can be applied to different domains where multidimensional temporal data exists. We have also designed novel 3D visualization methods to present and analyze large-scale data. Many researchers have likewise put much effort into high-level models of VA system design. University of British Columbia computer science professor Tamara Munzner has proposed a nested model for visualization design and validation with four layers: characterize the task and data in the vocabulary of the problem domain, abstract into operations and data types, design visual encoding and interaction techniques, and create algorithms to execute techniques efficiently.[15] Based on that model, UNC Charlotte researcher Xiaoyu Wang and his coauthors have proposed a two-stage framework for designing a VA system in organizational environments.[16] The design and evaluation of our system are inspired by these general design models and frameworks.

System Design

A special aspect of Apsara is its ability to access a computer's graphics processing unit (GPU). The GPU enables a computer to render complex 3D computer animations. In this section, we talk about the visualization techniques and interaction design in Apsara. We also give a high-level overview of the VA process.

Visualization Techniques

Visual Encodings

The basic idea of Apsara is to present large-scale data in a vast universe (see Figure 2). Each data point is represented by a standard sphere with different sizes in the 3D universe, enhanced by color, motion, and sound. The important attributes of multidimensional data can be mapped to these visual encodings with optional data models, thus allowing a compact overview of the whole data set. Apsara uses a time axis to visualize temporal data. In addition to displaying a static universe in a fixed time frame, it is able to show the dynamic evolution of the universe over any time period.

The basic idea of Apsara is to present large-scale data in a vast universe.

Customizable Visualization

One useful feature that is often ignored in visualization function design is customizable visualization.[17] Re-parameterized settings may lead to completely different data models and visual representations, even with the same data and visualization techniques. By changing visual parameter settings, users will gain different visual impressions and may have more opportunities to get insight into patterns, trends, and correlations in the data. In Apsara, customizable visualization is achieved by designing different data models for different user demands and use cases. Apsara also allows users to set their own workspace with a selected data set and build data models with customizable parameters.

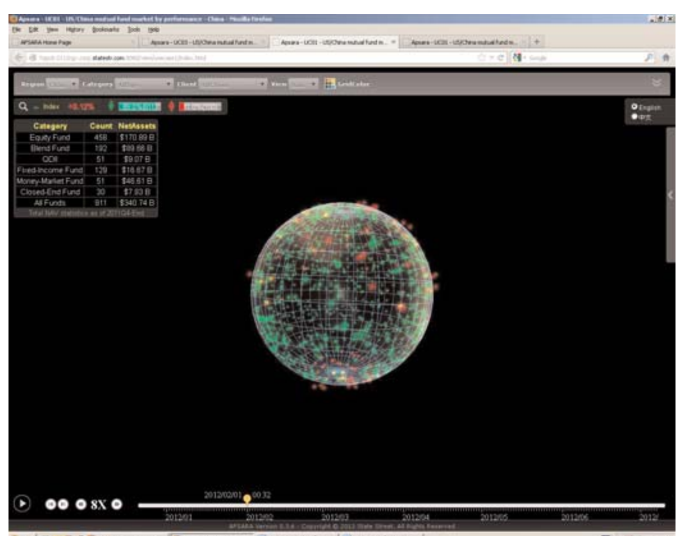

Figure 2 — The overview of the whole data set presented in a single universe. Each data point is represented by a standard sphere.

Interaction Design

Interaction is increasingly seen as central in representing individuals' analysis processes within VA systems. It is through the interactive manipulation of a visual interface that knowledge is constructed, tested, refined, and shared.[18] A considerable amount of information regarding a user's analysis process with a visualization tool is captured by interactions.[19] Apsara implements several interactive features:

- **Data filter.** By default, the universe presents the overview of the whole data set. Users may select certain types of data they are interested in, and the other data will be filtered from the universe.
- **Navigation.** Since the spheres are distributed in a 3D universe and there is an earth frame, it is unavoidable that spheres on the back side or inside the earth may be occluded by those on the front side or outside the earth. Apsara allows users to rotate the earth to explore it from different angles. Users can also zoom in to get a more detailed view of some special groups of spheres or spheres inside the earth.
- **Detail on demand.** When a sphere is far away from the universe center, it may have special attribute values that are greatly different from others. There may be great potential benefit to further investigating these "outliers." Users can click on the sphere and drill down for detailed information, which will be presented in a small window. Furthermore, we plan to add a HTTP link that will guide users to a website that offers background information and other external resources for such spheres.
- **Auto-play.** There is a time axis at the bottom of the universe. When the user drags the time axis, the universe will change accordingly to present the new state of the data set. However, it is difficult to present continuous change over a long time period by manual operation. So Apsara provides auto-play functionality to move the time axis automatically with optional speed-up, thus allowing users to put more focus on analysis of patterns, trends, and correlations.

Visual Analysis

The VA process combines automatic analysis and visualization methods with a tight coupling through human interaction in order to gain knowledge from data. Figure 3 shows an abstract overview of the stages and their transitions in the VA process.

In many application scenarios, heterogeneous data sources need to be integrated before visual or automatic analysis methods can be applied. This raw data often contains various errors and some of it is incomplete. Therefore, the first step is to pre-process and transform the data to derive different representations for further exploration. During the process, pre-analysis and verification are important for guaranteeing the quality of data. Other typical pre-processing tasks include data cleaning, normalization, grouping, and integration.

After the transformation, the characterization of data needs to be defined clearly so that the important attributes can be extracted out for visualization and analysis. Then the analyst may choose between applying visual or automatic analysis methods. If an automated analysis

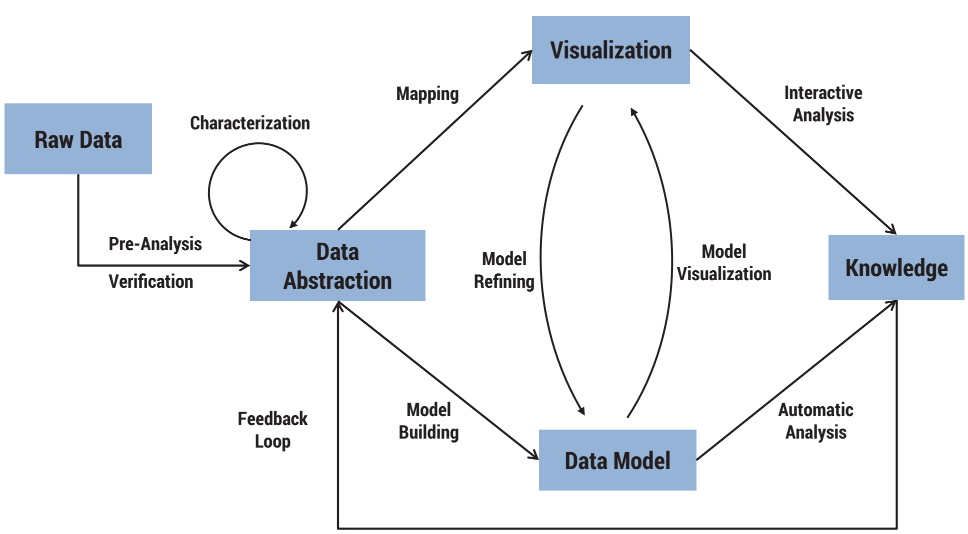

Figure 3 — The VA process in Apsara.

is used first, data mining methods are applied to build models of the data. Once a model is created, the analyst needs to evaluate and refine the model, which can best be done by interacting with the data. Visualizations allow the analyst to interact with the automatic methods by modifying parameters or selecting other analysis algorithms. Model visualization can then be used to evaluate the findings of the generated models. Next comes continuous refinement and verification of the preliminary model. Misleading results in an intermediate step can thus be discovered at an early stage, leading to more suitable data models and a higher confidence.

If visual data exploration is performed first, the user needs to confirm the generated hypotheses by interactive analysis. User interaction with the visualization can reveal insightful information; for instance, by zooming in on different data areas or by considering different visual views on the data. Findings in the visualizations can be used to steer model building in the automatic analysis. In summary, in the VA process, knowledge can be gained from visualization and automatic analysis, as well as from the preceding interactions between visualizations, models, and the human analysts.

Case Study: Visualizing China's Mutual Fund Market

In this section, we discuss the use of Apsara with real-world data from the Chinese mutual fund industry.

Background

In a 2012 report, Shanghai-based consulting firm Z-Ben Advisors projected that total assets under management (including public and non-public funds) in China's financial services industry would approach US $1 trillion by late 2015.[20] While such a high might not be permanently sustained, it provides a basis for additional investment into the industry on the part of managers. This will be buoyed by growth in non-public segments, itself driven by increases in demand from high net worth and institutional investors.

Highly unbalanced portfolios, both geographically and within asset classes, will add to the big change in China's financial services industry. The QDII (Qualified Domestic Institutional Investor) program, through which domestic investors can buy into global securities, offers considerable room for growth given its small current size. Global investors are also severely underinvested in China at present, and the rebalancing of portfolios, as well as the emergence of the Greater China asset class (separate from BRICS [Brazil, Russia, India, and China]), will provide significant impetus for such a transition.

There are currently several large pools of assets on the Chinese mainland that have not yet been put to work. These range from insurers' own capital to pension and other centralized funds that are highly concentrated into cash and fixed income. These imbalances alone do not mean that investments will be reallocated. Rather, the gap serves as the basis for assumptions on how asset flows will be redirected over the coming years in order to seize potential opportunities for entrants to or current participants in China's financial services industry. In addition to greatly impacting the financial services industry, growth will promote small and medium-sized private enterprises, which today are significantly underdeveloped within the domestic fund industry. Traditionally, banks in China tend to offer loans to large state-owned companies instead of private companies. Thus, most private companies can only get money in abnormal ways, with a very high lending rate. Growth in financial services offerings will therefore become even more important as such firms face rising financing costs.

Traditional analysis with "basic" visualization methods cannot offer a compact overview of the whole mutual fund market, and other analysis tools powered by "slightly advanced" visualization methods lack detailed information about single funds or certain groups of mutual funds.

As we look at China's mutual fund industry, we need a powerful way to explore and demonstrate this whole process and help decision making as well as policy making. Traditional analysis with "basic" visualization methods cannot offer a compact overview of the whole mutual fund market, and other analysis tools powered by "slightly advanced" visualization methods lack detailed information about single funds or certain groups of mutual funds. But visual analytics, which combines both automatic, intelligent data analysis methods with highly effective visualization and interaction capabilities,[21] can turn large-scale complex data into reliable and comprehensible knowledge.

Briefly, our goal is to find ways to:

- Visualize the whole mutual fund market while retaining the information on a single mutual fund as specifically as possible
- Visualize the market's dynamic changes over periods of time
- Support domain analysts' work

The first and most crucial step in designing a visual analytics system is to define the goal and characterization of the data, which will determine the choice of visual encodings and the data model.

Data Characterization

The first and most crucial step in designing a VA system is to define the goal and characterization of the data,[22] which will determine the choice of visual encodings and the data model. A mutual fund is a type of professionally managed collective investment vehicle that pools money from many investors to purchase securities. There is a wide variation in mutual fund data, which is always high-dimensional and changes dynamically. The characteristics of mutual funds examined in Apsara are:

- **Net asset value (NAV).** This is the value of an entity's assets less the value of its liabilities, often in relation to open-end or mutual funds, since shares of such funds registered with the US Securities and Exchange Commission are redeemed at their net asset value.
- **Growth rate and volatility.** These both measure the performance of a mutual fund. A negative growth rate is usually considered bad. A mutual fund with high volatility often means high return with high risk, which makes it difficult to judge.
- **Fund type.** Since "mutual fund" is a broad concept, all mutual fund data is classified into one of seven types according to a fund's principal investments: equity fund, blend fund, QDII fund, fixed-income fund, short-term debt fund, money market fund, and closed-end fund. As noted above, QDII is a program relating to the capital market established to allow financial institutions to invest in foreign securities and bonds. In this article, it mainly refers to China QDII.
- **Client.** This is the financial company that runs a certain mutual fund.

Analysis of the NAV of all Chinese mutual funds revealed a very large gap between the different funds. The NAV ratio can be nearly 1,000:1 when comparing the largest and smallest mutual fund. This is a very important characteristic of the Chinese mutual fund market that should be taken into consideration when building the data model.

Data Management

Data management, which aims to ensure data consistency, avoid duplication, and handle data transactions, is important for VA tools. Most of the data in Apsara is fetched from different websites where the format and content of data vary a lot. Before applying visualization techniques to build the data model, this heterogeneous data must be pre-processed, integrated, and verified.

Initially, the data is fetched from two data sources (CSRC, which is about Chinese fund data, and Morningstar, which is about US fund data) with Apache HttpClient, an open source Web crawling tool. The data contains fund name, fund type, fund code, establishing date, and fund value for a whole year. (Our data was from 2012.)

It's necessary to deal with the data quality problem after collecting the data, as the data could be incomplete, inconsistent, or contain measurement errors. An "experts voting" strategy is adopted to tackle this problem. That is, another four powerful data sources are added to complement the data. When analysts meet with missing or inconsistent data, it is replaced with the most reliable value from these four data sources by weighted voting.

Basic Data Model

This section describes in detail about how to choose the visual encodings and how to build the basic data model so that it is as intuitive as possible.

Visualizing Net Asset Value

Sphere volume is a reasonable and intuitive way to represent a mutual fund's NAV. However, the characteristics of real data described above suggest that NAV cannot be mapped to volume directly. Small spheres will be occluded by several extremely large spheres, thus destroying the delicate sense of the whole market. So the Log function is assigned to the NAV of each mutual fund as a small adjustment. With this adjustment, small and large spheres can still be told apart clearly, while the whole universe looks more in harmony.

Nevertheless, it is difficult to compare the volume of two spheres that are both small or both big when they are distributed in three-dimensional space. What's more, one might wonder how to "distribute" these spheres and what is the meaning of their position in the universe? In order to solve these problems, we added several features to Apsara. First, we added a center point of the universe. The distance from a sphere to the universe center is based on the volume (i.e., NAV of the mutual fund) of the sphere. Larger spheres are further from the universe center. However, if NAV is mapped to the distance directly, many small spheres will stay close to the universe center, while the largest sphere will be too far away to be shown on the screen (the largest ratio is about 1000:1). The broad space between the small and large spheres will be empty, making the universe unacceptable. Our solution was to make use of the basic formula that calculates volume of a sphere:

$$V = \frac{4}{3}\pi R^3$$

Here V is the volume of a sphere and also the NAV of a mutual fund. R is the radius of this sphere and also the distance from the mutual fund to the universe center. So another equation about the distance R can be derived:

$$R = \sqrt[3]{\frac{3V}{4\pi}}$$

By using this equation to calculate the distance between a sphere and the universe center, the largest ratio will be reduced to about 9:1. With this transformation, the overall distribution is more reasonable. Now the spheres can be divided into different groups according to their distance to the universe center. Spheres in the same group have the same distance to the universe center and are distributed on the surface of the same sphere, which can be regarded as a virtual sphere.

Visual Mapping of Fund Type and Client

The spheres cannot be distributed randomly on the surface since other characteristics of mutual fund data also need to be visualized. People are familiar with longitude and latitude when talking about spheres. So we take advantage of them as the new visual encodings. As noted earlier, all mutual funds data are classified into six types by their principal investments. Now each type of mutual fund is mapped to different latitude intervals according to its proportion in the market. The longitude represents the client of each mutual fund.

In order to give users an intuitive sense of a "virtual index," a new sphere is added in the universe. The center is the center of the universe, and the volume of this special index sphere is also meaningful. It represents the average NAV of the whole mutual fund market so that the radius of this index has the same aspect ratio as the spheres. The index acts as a benchmark, with spheres outside the earth having a higher NAV than the average value of the market. Now users can easily compare the NAV of different mutual funds by checking their distance to the universe center and get a general idea of NAV size by checking the positional relation between the index and the center.

You might have noticed that the distance from a sphere to the universe center and the volume of the sphere both represent a mutual fund's NAV. This visualization method is called "redundant encoding," which is recommended to help increase accuracy and enhance perception.

Visualizing Growth Rate and Volatility

Finally, the growth rate and volatility of mutual funds should be visualized in the universe. Since mutual fund data is time series data, both growth rate and volatility are highly related to time. In order to visualize the change of the mutual fund market over time, a time axis is added at the bottom of the universe. When users drag the time axis, a sphere (representing a single fund) will move according to the change of its NAV. Say, for example, that in 2011 a fund has a higher NAV than the average value of the whole market. In that case, it will stay outside the virtual earth. If the fund's NAV shrinks below the average value in 2012, then it will move inside the virtual earth. Thus, users can get a general idea of a fund's growth rate and volatility by checking the movement speed and range of its sphere.

Traditionally, mutual funds are compared with a stock market index, such as the Dow Jones Industrial Average in the US and the Shanghai Composite Index in China. If users want to know more about a fund's growth rate, they can check the sphere's color. In Apsara, when visualizing China's mutual fund data, red spheres outperform the market, while green spheres underperform the market.

2D vs. 3D

This Apsara data model is aimed at common characteristics of a single mutual fund — data that can be visualized via traditional 2D charts as well. However, it is innovative for Apsara to be able to present all mutual funds data on the same screen (i.e., the 3D universe), thereby offering a compact overview of the whole mutual fund market while retaining information on individual mutual funds as specifically as possible.

Customizable Visualization

The data model described in the previous section mainly focuses on the visualization of fund NAV. It is difficult to have an idea of the growth rate compared to a certain day in the past. So another data model was designed to deeply explore the growth rate of a mutual fund (see Figure 4). In this model, the distance from a sphere to the universe center is calculated by the equation below:

$$R = (1 + c) * \alpha, c \in [-1, +\infty)$$

Here R is the distance and c is the growth rate of a mutual fund compared to a certain day in the past. Since the growth rate can be negative — whereas negative distance cannot be visualized in the universe — the model adds a 1 to all growth rate values to guarantee that the value is always non-negative. Then the growth rate is multiplied by a constant number α (default value is 200) to expand the distance so that the universe looks more reasonable.

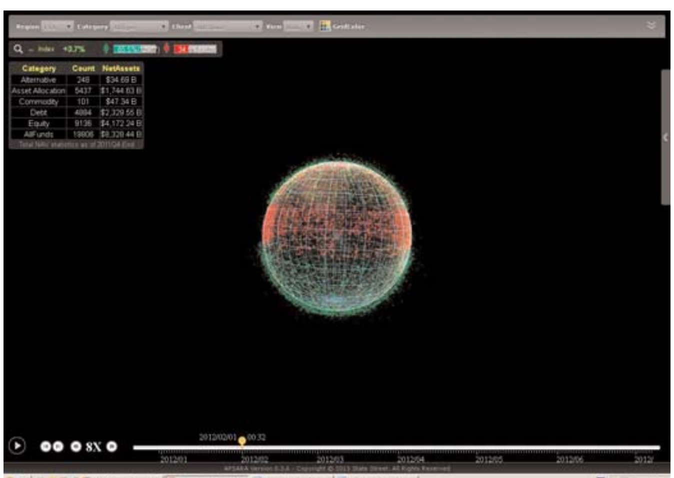

Figure 4 — Another data model to visualize the growth rate of mutual funds in the US market. While the index is growing up, most spheres with low latitude are red (which, in the US context, denotes underperforming the market). Through interactive analysis, we found that the "debt fund" accounted for the majority of the red spheres. This shows that the debt fund is less affected by the stock market.

In this scenario, the benchmark (i.e., the virtual earth) no longer represents the average NAV of the market. Now it represents the growth rate of the stock market index (e.g., the Dow Jones) compared to a certain day in the past. The new equation to calculate the earth radius is the same as for the mutual fund spheres. In this way, users can easily compare mutual funds' performance with the stock market index by checking the position of the spheres. Spheres outside the earth outperform the market, while spheres inside the earth underperform the market. The colors green and red are also used to enhance perception, as described earlier.

Application Examples

In this section, we present two use cases in Apsara where real data from both the China and US mutual fund markets is applied. The first case shows the evolution of the Chinese mutual fund market in 2012 using interactive operations. The second case demonstrates Apsara's capability in forecasting the development of China QDII.

Visualizing the Mutual Fund Market in China and the US

In 2012 (the year for which we had data), there were 21,686 mutual funds in the US and 869 funds in China. While the US has far more funds than China, the difference of earth radius (when representing average NAV) in Apsara is quite small since it is adjusted by the equation described above. But it is still intuitive that the US market is more prosperous (see Figure 5). The virtual earth for the US is almost completely covered by spheres representing funds. In the China market, parts of the earth's surface (such as the space for QDII) are nearly empty.

We described above how the real mutual fund data is fetched from different data sources. After being preprocessed and integrated, this real data is applied to Apsara. Interactive operations are also fully examined in this case. By dragging the time axis or clicking the auto-play button, users can see the evolution of the mutual fund market. As time goes by, spheres will change their color and move in or out of the earth. A few spheres will vanish when certain mutual funds exit the market, and new spheres will also come into the universe.

After using the data filter to filter out mutual funds that do not belong to QDII, it is clear that QDII in China is still in the early stage of growth. Several QDII "outliers" are also found to have much higher NAV than other QDII spheres. Through further investigation of their companies' background, the potential business benefit and investment chances are confirmed.

Predictive Analysis of China QDII

At the beginning of 2012, the total value of QDII in China was just 9.07 billion RMB, accounting for only 2.5% of all mutual fund market assets (340.74 billion). While China QDII is just "breaking the ice," it is supposed to have great potential benefit for custodian businesses. (A custodian bank is a special financial institution that is responsible for protecting the financial assets of individuals or companies against the fund managers.) Since China QDII is for Chinese financial institutions to invest in foreign countries, there will be great demand for global custodians while a lot of China QDII funds are set up.

Apsara enables analysts to make predictions about the mutual fund market. The goal is not to predict the growth rate with pinpoint accuracy, but just to show the possible trends in the mutual fund market. In this case, the predictive model in Apsara is based on Z-Ben Advisors' internal research work.

According to that predictive model, in 2018 the total NAV of China QDII will have grown to 53.69 billion RMB, about six times larger than the 2012 figure. The worldwide mutual fund asset value will be growing more slowly than China QDII. Since the capital of China QDII is increasing fast, when an analyst drags the time axis to a future time, many new spheres of QDII will be born in the universe (see Figures 6 and 7). Auto-playing the market change from 2012 to 2018 clearly shows the trend of QDII in the future, thus helping financial analysts to make better decisions and convince officers to adopt new policies.

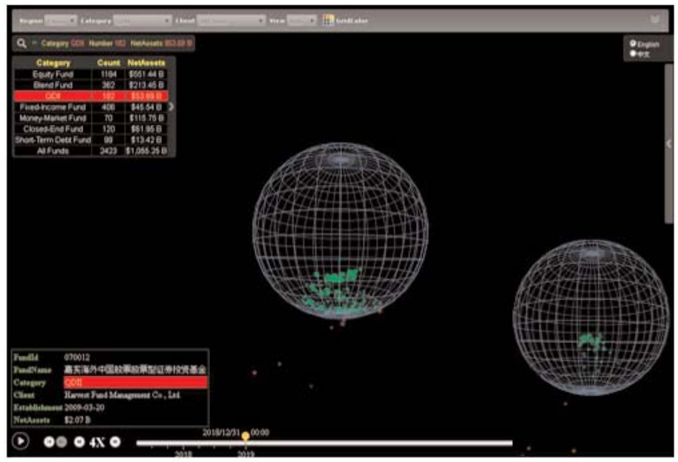

Figure 6 — China QDII in 2018 (projected) with a comparison view. Other types of funds are all filtered out. The lower righthand corner shows the universe in 2012. QDII's rapid growth can be intuitively understood.

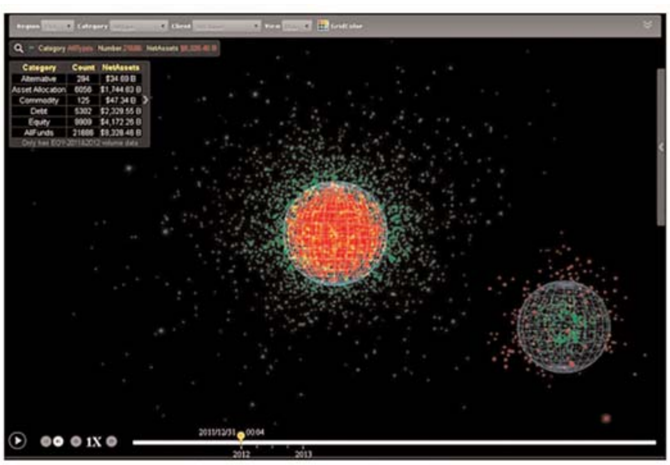

Figure 5 — Overview of the US mutual fund market in 2012. There are 21,686 spheres in total. The lower righthand corner shows the universe of the China mutual fund market for comparison.

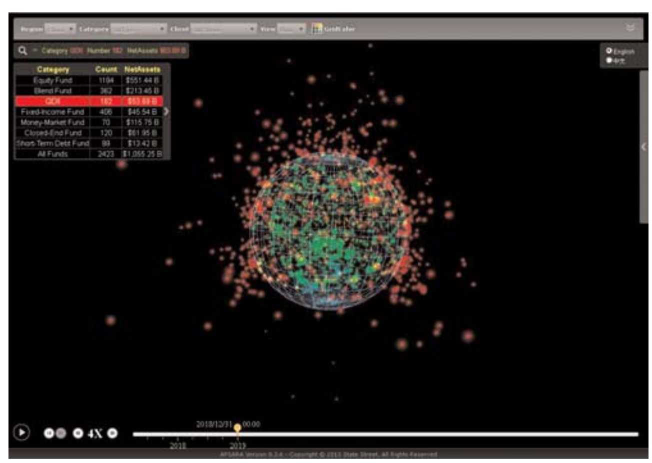

Figure 7 — Predictive result of China QDII in 2018 with all mutual funds shown in the universe. Spheres that represent QDII are highlighted in blue. This visualization reveals that both the total NAV and quantity of QDII have great potential to increase.

Conclusion

The best way to visualize large-scale complicated data remains an open question. While traditional 2D charts are good at presenting analysis results precisely and concisely, they are limited when it comes to supporting information exploration. We believe that a compact overview of an entire data set with handy human-computer interactions would greatly help in detecting patterns and insights. Thus, we developed a 3D VA system for multidimensional temporal data that enables users to intuitively explore and gain insights from raw data. This customizable solution can be adapted to different data models and applied to multiple domains.

Acknowledgments

We are indebted to Prof. Jianling Sun, Prof. Shanping Li, Prof. Bo Zhou, Prof. Bin Xu, and Ms. Jin Xu for their advice and assistance with this article. The design and implementation of the visual analytics prototype was the work of research teams at State Street Hangzhou and Zhejiang University, including Yangfan Chen, Dawei Mei, Haifeng Huang, Jie Mou, Junfang Cheng, Kang Yang, Rongrong Chen, Xuanyu Zhao, Yuchen Wu, Cun Liu, Fangzhao Xing, Zhenlei Qian, Minmin Lin, Li Li, Chenghao Liu, and Lihan Ju. This work was funded by State Street Technology (Zhejiang), a wholly owned subsidiary of State Street Corporation.

Endnotes

[1]Keim, Daniel A., et al. "Challenges in Visual Data Analysis." *Proceedings of the Tenth International Conference on Information Visualization*. IEEE, 2006 (http://ieeexplore.ieee.org/document/1648235/).

[2]Keim, Daniel A., et al. (eds.) *Mastering the Information Age — Solving Problems with Visual Analytics*. Eurographics Association, 2010 (www.vismaster.eu/wp-content/uploads/2010/11/VisMaster-book-lowres.pdf).

[3]Keim, Daniel A., et al. "Visual Analytics: Scope and Challenges." In *Visual Data Mining*, edited by Simeon J. Simoff, Michael H. Böhlen, and Arturas Mazeika. Springer-Verlag Berlin Heidelberg, 2008 (https://link.springer.com/chapter/10.1007/978-3-540-71080-6_6).

[4]Keim et al. (see 2).

[5]Tufte, Edward R. *The Visual Display of Quantitative Information*. Vol. 2. Graphics Press, 1983.

[6]Keim et al. (see 1).

[7]Tufte (see 5).

[8]Zhang, Leishi, et al. "Visual Analytics for the Big Data Era — A Comparative Review of State-of-the-Art Commercial Systems." *Proceedings of the IEEE Conference on Visual Analytics Science and Technology 2012*. IEEE, 2012 (http://ieeexplore.ieee.org/document/6400554/).

[9]Harger, John R., and Patricia J. Crossno. "Comparison of Open-Source Visual Analytics Toolkits." *Proceedings of the SPIE Conference on Visualization and Data Analysis 2012*. SPIE, 2012 (http://proceedings.spiedigitallibrary.org/proceeding.aspx?articleid=1283685).

[10]Shen, Zeqian, and Kwan-Liu Ma. "MobiVis: A Visualization System for Exploring Mobile Data." *Proceedings of the IEEE Pacific Visualization Symposium*. IEEE, 2008 (http://ieeexplore.ieee.org/document/4475474/).

[11]Guo, Diansheng, et al. "A Visualization System for Space-Time and Multivariate Patterns (VIS-STAMP)." *IEEE Transactions on Visualization and Computer Graphics*, Vol. 12, No. 6, November-December 2006 (http://ieeexplore.ieee.org/document/1703367/).

[12]Chang, Remco, et al. "WireVis: Visualization of Categorical, Time-Varying Data from Financial Transactions." *Proceedings of the IEEE Symposium on Visual Analytics Science and Technology 2007*. IEEE, 2007 (http://ieeexplore.ieee.org/document/4389009/).

[13]Xu, Weijia, et al. "Analysis of Large Digital Collections with Interactive Visualization." *Proceedings of the IEEE Conference on Visual Analytics Science and Technology 2011*. IEEE, 2011 (http://ieeexplore.ieee.org/document/6102462/).

[14]Fisher, Danyel. "Hotmap: Looking at Geographic Attention." *IEEE Transactions on Visualization and Computer Graphics*, Vol. 13, No. 6, November-December 2007 (http://ieeexplore.ieee.org/document/4376139/).

[15]Munzner, Tamara. "A Nested Model for Visualization Design and Validation." *IEEE Transactions on Visualization and Computer Graphics*, Vol. 15, No. 6, November-December 2009 (http://ieeexplore.ieee.org/document/5290695/).

[16]Wang, Xiaoyu, et al. "A Two-Stage Framework for Designing Visual Analytics System in Organizational Environments." *Proceedings of the IEEE Conference on Visual Analytics Science and Technology 2011*. IEEE, 2011 (http://ieeexplore.ieee.org/document/6102463/).

[17]Zhang et al. (see 8).

[18]Pike, William A., et al. "The Science of Interaction." *Information Visualization*, Vol. 8, No. 4, January 2009 (http://journals.sagepub.com/doi/abs/10.1057/ivs.2009.22).

[19]Wang et al. (see 16).

[20]Z-Ben Advisors. "China: The World's Best Opportunity for Asset Managers?" Citibank, 2012 (www.citibank.com/transactionservices/home/about_us/articles/docs/china_asset.pdf).

[21]Schreck, Tobias, et al. "Trajectory-Based Visual Analysis of Large Financial Time Series Data." *ACM SIGKDD Explorations Newsletter*, Vol. 9, No. 2, December 2007 (http://dl.acm.org/citation.cfm?id=1345454).

[22]Munzner (see 15).

Jerry Cristoforo is Executive VP and CTO for State Street Corporation. In addition, Mr. Cristoforo is Chairman and General Manager of State Street Technology (Zhejiang), where he pioneered State Street's partnership with Zhejiang University in Hangzhou, China, in 2000. He is a senior computer/securities industry executive with in-depth technology and management experience in advanced distributed systems architecture, product development, workflow reengineering, large-scale systems integration, and project management. He can be reached at jacristoforo@statestreet.com.

Qiao Huang is a PhD student in the College of Computer Science and Technology at Zhejiang University. His research interests include mining software repositories and empirical software engineering. He can be reached at tkdsheep@zju.edu.cn.

Zhiyu Peng is Principal Engineer at Insigma Hengtian, a joint venture of State Street Corporation and Insigma Corporation in China. Dr. Peng has been driving the innovation activities and incubating R&D topics in Hengtian, and his research focuses on big data, data mining, and data visualization. He can be reached at zhiyupeng@hengtiansoft.com.

Xiaohu Yang is Executive Vice Dean and Professor of the College of Computer Science and Technology at Zhejiang University. Dr. Yang is the cofounder of State Street Zhejiang University Technology Centre, a joint research center set up in 2001 between State Street Corporation and Zhejiang University for advanced research and development of global financial software systems and technologies. Dr. Yang's research focus includes software engineering, fintech, and cloud computing. He can be reached at yangxh@zju.edu.cn.

Digital Transformation & Innovation Bootcamp

Karim Lakhani
Professor, Harvard Business School
Fellow, Cutter Consortium

Next Bootcamp:

October 25-27

Babson Executive Conference Center
Wellesley, MA USA

The Digital Transformation & Innovation Bootcamp program fee is US $15,000 for a team of 3–5, plus an additional fee per person for accomodations and materials.

For more information or to register your team, contact your Cutter Consortium Account Executive at +1 781 648 8700 or sales@cutter.com.

Is Your Business Model Still Relevant? Will It Be Next Year?

The competitive landscape is nothing like we've seen before. An organization in a completely different industry may be poised to swallow up your market — and you might not even see it coming.

Did Nokia — a cell phone company — see Apple — a niche personal computer company — sneaking in to pull the market out from under it?

There's much to learn from Nokia's and others' experiences about what to do, what not to do, how to plan for, and how to manage a shift toward creating and capturing new value. At Cutter Consortium's *Digital Transformation & Innovation Bootcamp* you'll dive deep into what you need to do to succeed.

Designed for just a handful of teams from noncompeting firms, the *Bootcamp* provides a rigorous process for you to revolutionize your digital transformation journey.

The Agenda

Case Studies & Lectures. Harvard Business School Professor and Cutter Fellow Karim Lakhani will lead you through several case studies and lectures. Prior to each case discussion, you'll work through its issues from the variety of perspectives you enjoy when collaborating with colleagues from outside your industry.

Hands-on Working Sessions. Based on the framework you'll learn, you'll begin working on a transformation project that will power your organization. During multiple Transformation Sprints, your team of 3-5 colleagues from your own organization will develop new ideas and new business models to help your business thrive. You'll be ready to prototype the resulting new/revised project or strategy when you get back home.

Project Assessment. You'll get feedback on your transformation project from the other teams. The fresh perspective of executives from other sectors/noncompeting companies will uncover insights about your situation that you'd otherwise likely miss.

Register Your Team Now

In just 2 days, Professor Lakhani will guide you through an examination of how digital innovation is transforming our business landscape. You'll learn:

- What's truly unique about the new digital strategies that are changing our environment?
- What digital strategies have established organizations employed?
- What strategies might be right for your organization?

With only one organization per industry permitted, spaces are filling quickly!

For more information, or to reserve your team's spot for our October 25-27 *Digital Transformation & Innovation Bootcamp,* call us now at +1 781 648 8700 or email sales@cutter.com.

Cutter Consortium
Access to the Experts

Cutter Business Technology JOURNAL

About Cutter Consortium

Cutter Consortium is a unique, global business technology advisory firm dedicated to helping organizations leverage emerging technologies and the latest business management thinking to achieve competitive advantage and mission success. Through its research, training, executive education, and consulting, Cutter enables digital transformation.

Cutter helps clients address the spectrum of challenges technology change brings — from disruption of business models and the sustainable innovation, change management, and leadership a new order demands, to the creation, implementation, and optimization of software and systems that power newly holistic enterprise and business unit strategies.

Cutter pushes the thinking in the field by fostering debate and collaboration among its global community of thought leaders. Coupled with its famously objective "no ties to vendors" policy, Cutter's *Access to the Experts* approach delivers cutting-edge, objective information and innovative solutions to its clients worldwide.

For more information, visit www.cutter.com or call us at +1 781 648 8700.

CUTTER CONSORTIUM
Access to the Experts