


LEARNING MADE EASY



Tangent Works
Special Edition

Predictive Analytics for Time Series with InstantML

for
dummies[®]
A Wiley Brand



Obtain business value
from time-series data

Automate forecasting and
anomaly detection

Democratize machine
learning

Brought to
you by:



Tangent Works
TIM

Lawrence Miller

About Tangent Works

Tangent Works is a machine learning company specializing in time-series data. TIM, its Tangent Information Modeler, is an automated predictive model-building engine that creates human-readable models from historical time-series data. These models provide users with high-quality forecasts and anomaly detections, creating business value in use cases across industries.

Predictive Analytics for Time Series with InstantML

**for
dummies[®]**
A Wiley Brand



Predictive Analytics for Time Series with InstantML

Tangent Works Special Edition

by Lawrence Miller

for
dummies[®]
A Wiley Brand

Predictive Analytics for Time Series with InstantML For Dummies®, Tangent Works Special Edition

Published by

John Wiley & Sons, Inc.

111 River St.

Hoboken, NJ 07030-5774

www.wiley.com

Copyright © 2021 by John Wiley & Sons, Inc., Hoboken, New Jersey

No part of this publication may be reproduced, stored in a retrieval system or transmitted in any form or by any means, electronic, mechanical, photocopying, recording, scanning or otherwise, except as permitted under Sections 107 or 108 of the 1976 United States Copyright Act, without the prior written permission of the Publisher. Requests to the Publisher for permission should be addressed to the Permissions Department, John Wiley & Sons, Inc., 111 River Street, Hoboken, NJ 07030, (201) 748-6011, fax (201) 748-6008, or online at <http://www.wiley.com/go/permissions>.

Trademarks: Wiley, For Dummies, the Dummies Man logo, The Dummies Way, Dummies.com, Making Everything Easier, and related trade dress are trademarks or registered trademarks of John Wiley & Sons, Inc. and/or its affiliates in the United States and other countries, and may not be used without written permission. All other trademarks are the property of their respective owners. John Wiley & Sons, Inc., is not associated with any product or vendor mentioned in this book.

LIMIT OF LIABILITY/DISCLAIMER OF WARRANTY: THE PUBLISHER AND THE AUTHOR MAKE NO REPRESENTATIONS OR WARRANTIES WITH RESPECT TO THE ACCURACY OR COMPLETENESS OF THE CONTENTS OF THIS WORK AND SPECIFICALLY DISCLAIM ALL WARRANTIES, INCLUDING WITHOUT LIMITATION WARRANTIES OF FITNESS FOR A PARTICULAR PURPOSE. NO WARRANTY MAY BE CREATED OR EXTENDED BY SALES OR PROMOTIONAL MATERIALS. THE ADVICE AND STRATEGIES CONTAINED HEREIN MAY NOT BE SUITABLE FOR EVERY SITUATION. THIS WORK IS SOLD WITH THE UNDERSTANDING THAT THE PUBLISHER IS NOT ENGAGED IN RENDERING LEGAL, ACCOUNTING, OR OTHER PROFESSIONAL SERVICES. IF PROFESSIONAL ASSISTANCE IS REQUIRED, THE SERVICES OF A COMPETENT PROFESSIONAL PERSON SHOULD BE SOUGHT. NEITHER THE PUBLISHER NOR THE AUTHOR SHALL BE LIABLE FOR DAMAGES ARISING HEREFROM. THE FACT THAT AN ORGANIZATION OR WEBSITE IS REFERRED TO IN THIS WORK AS A CITATION AND/OR A POTENTIAL SOURCE OF FURTHER INFORMATION DOES NOT MEAN THAT THE AUTHOR OR THE PUBLISHER ENDORSES THE INFORMATION THE ORGANIZATION OR WEBSITE MAY PROVIDE OR RECOMMENDATIONS IT MAY MAKE. FURTHER, READERS SHOULD BE AWARE THAT INTERNET WEBSITES LISTED IN THIS WORK MAY HAVE CHANGED OR DISAPPEARED BETWEEN WHEN THIS WORK WAS WRITTEN AND WHEN IT IS READ.

For general information on our other products and services, or how to create a custom *For Dummies* book for your business or organization, please contact our Business Development Department in the U.S. at 877-409-4177, contact info@dummies.biz, or visit www.wiley.com/go/custompub. For information about licensing the *For Dummies* brand for products or services, contact BrandedRights&Licenses@Wiley.com.

ISBN 978-1-119-76753-4 (pbk); ISBN 978-1-119-76755-8 (ebk)

Manufactured in the United States of America

10 9 8 7 6 5 4 3 2 1

Contents at a Glance

Introduction	1
CHAPTER 1: Recognizing the Business Value of Predictive Analytics on Time-Series Data	5
CHAPTER 2: Understanding the Challenges of Implementing Predictive Analytics.....	13
CHAPTER 3: Exploring Business Use Cases	21
CHAPTER 4: Why Time-Series Forecasting Is Difficult.....	35
CHAPTER 5: Exploring Time-Series Predictive Analytics	43
CHAPTER 6: Implementing the Tangent Information Modeler and Instant Machine Learning.....	51
CHAPTER 7: Using the Tangent Information Modeler on Other Platforms	57
CHAPTER 8: Ten Ways to Get Value from the Tangent Information Modeler	65

Publisher's Acknowledgments

We're proud of this book and of the people who worked on it. Some of the people who helped bring this book to market include the following:

Project Editor: Elizabeth Kuball

Acquisitions Editor: Ashley Coffey

Editorial Manager: Rev Mengle

Business Development

Representative: Molly Daugherty

Production Editor:

Mohammed Zafar Ali

Special Help:

Elke Van Santvliet,

Ján Dolinsky,

Henk De Metsenaere,

Scott Bergquist, Dirk Michiels

Table of Contents

INTRODUCTION	1
About This Book	1
Foolish Assumptions	2
Icons Used in This Book	3
Beyond the Book	3
CHAPTER 1: Recognizing the Business Value of Predictive Analytics on Time-Series Data	5
Creating Business Value through Predictive Analytics	5
Understanding the Value Proposition of Data and Analytics	8
Defining a Business Value Model	9
CHAPTER 2: Understanding the Challenges of Implementing Predictive Analytics	13
Data Quality and Availability	14
Project and Operation Governance	15
Lack of a Data Strategy	17
Being Stuck in Handcrafted Modeling and Failing to Move to Production	17
Structural Changes in the Environment Rendering Existing Predictive Models Obsolete	18
CHAPTER 3: Exploring Business Use Cases	21
Retail and Consumer Packaged Goods	21
Sales forecasting	22
Demand planning	23
Energy Sector	24
Energy consumption	25
Wind production	26
Solar production	26
Finance	27
Mortgage prepayment and credit	28
Manufacturing and Supply Chain Management	28
Predictive maintenance and asset health monitoring	29
Asset failure analysis	30
Supply chain strategy planning	30
Telecom	31
Testing customer loyalty with experience issues	32
5G: The next telco battleground	32

CHAPTER 4:	Why Time-Series Forecasting Is Difficult	35
	Circumstances Change Over Time	35
	Structural Changes	36
	Data Availability	38
	Need for Multiple Forecasts over Multiple Time Spans	39
CHAPTER 5:	Exploring Time-Series Predictive Analytics	43
	How to Build a Predictive Analytics Solution	43
	Handcrafted Models	45
	AutoML	46
	A New Paradigm: InstantML	47
CHAPTER 6:	Implementing the Tangent Information Modeler and Instant Machine Learning	51
	The Tangent Information Modeler Engine	51
	The Tangent Information Modeler Application	
	Programming Interface	52
	TIM Studio	53
CHAPTER 7:	Using the Tangent Information Modeler on Other Platforms	57
	Cloud Platforms	59
	Analytics and Business Intelligence Platforms	59
	Data Integration Platforms	60
	Machine Learning Platforms	62
	Internet of Things Platforms	63
CHAPTER 8:	Ten Ways to Get Value from the Tangent Information Modeler	65
	Driving Digital Transformation	65
	Focusing on Business Value	65
	Going Beyond Experimentation	66
	Getting Business Value through Augmented Analytics	66
	Approaching Time-Series ML in a New Way with InstantML	66
	Getting Results Fast	67
	Automating the Model-Building Process	67
	Generating Accurate Models	67
	Explaining Model Insights	68
	Integrating in Your Existing Landscape Easily	68

Introduction

Modern businesses are under constant pressure to reduce time to value. The key to success lies in valuable insights that are often locked away in massive amounts of data. This data may be collected in real time for specific purposes, but it can become stale if it isn't quickly analyzed and interpreted. The sources of data are now practically limitless, and the amount of data has become so massive that manually analyzing it is no longer possible. Machine learning (ML) has the potential to surface immense value in this data, but ML models are often stuck in the experimental phase being trained and retrained by expert data scientists and never used by business.

Machine learning is a popular buzzword in technology today, but handcrafted model building by experts isn't agile enough to come up with solutions that drive real-time business decisions and operations. Automated machine learning (AutoML) partially automates the model-building process to overcome some of these challenges and increase agility. Today, advancements in ML are creating new opportunities to further automate and accelerate the model-building process.

Some of these new opportunities are found in the field of time-series analysis. Time-series data is a sequence of data ordered over time. A popular form of analysis for time-series data is forecasting. Apart from forecasting, many applications can be found for anomaly detection. In fact, the potential use cases for forecasting and anomaly detection on time-series data are practically endless. Due to the unique challenges time-series data presents, AutoML isn't agile enough for time-series problems. InstantML is a new approach to model building that focuses solely on time-series data and overcomes the challenges of time-series problems.

About This Book

Predictive Analytics for Time Series with InstantML For Dummies consists of eight chapters that explore the following subjects:

- » Predictive analytics on time-series data from the business perspective (Chapters 1–3)

- » Predictive analytics on time-series data from the data-science perspective (Chapters 4–5)
- » The Tangent Information Modeler, or TIM (Chapters 6–8)

Each chapter is written to stand on its own, so if you see a topic that piques your interest, feel free to jump ahead to that chapter. You can read this book in any order that suits you (though I don't recommend upside down or backward).

Foolish Assumptions

It's been said that most assumptions have outlived their usefulness, but I assume a few things nonetheless:

- » Mainly, I assume that you work with lots of data and need to derive business value from that data.
- » Perhaps you're a business user who wants to learn how to implement predictive analytics, or you're a data scientist (or citizen data scientist) who wants to understand the value of augmented predictive analytics.
- » Perhaps you work in an IT department and need to learn more about how to extend your existing solutions with predictive analytics.

This book is written primarily for nontechnical readers who don't necessarily know a lot about the underlying technologies such as ML and artificial intelligence (AI).

If any of these assumptions describes you, then this is the book for you. If none of these assumptions describes you, keep reading anyway — it's a great book, and you'll learn quite a bit about predictive analytics.

Icons Used in This Book

Throughout this book, I use special icons in the margin to call attention to important information. Here's what to expect:



REMEMBER

The Remember icon points out important information you should commit to your nonvolatile memory, your gray matter, or your noggin — along with anniversaries and birthdays!



TECHNICAL
STUFF

If you seek to attain the seventh level of NERD-vana, perk up! This icon explains the jargon beneath the jargon and is the stuff nerds are made of!



TIP

Tips are appreciated, never expected — and I sure hope you'll appreciate these useful nuggets of information.

Beyond the Book

There's only so much I can cover in this short book, so if you find yourself at the end, thinking, "Where can I learn more?," check out <https://tangent.works>.

IN THIS CHAPTER

- » Understanding the business value of predictive analytics
- » Delivering value throughout your organization
- » Discovering how machine learning delivers more and faster business value

Chapter **1**

Recognizing the Business Value of Predictive Analytics on Time-Series Data

In this chapter, you explore how predictive analytics creates value for businesses across different industries, how to define the value proposition of predictive analytics, and the challenges that are unique to time-series forecasting and anomaly detection.

Creating Business Value through Predictive Analytics

Businesses have moved well beyond wondering what happened or why something happened. To be successful in today's hypercompetitive global market, they need to understand what *will* happen and what specific actions *will* drive the best outcomes. This is

the business value of predictive (and prescriptive) analytics (see Figure 1-1).

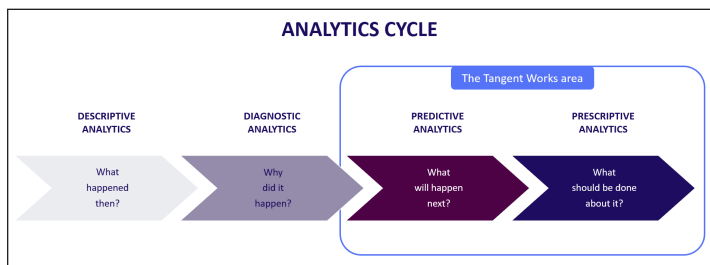


FIGURE 1-1: Predictive and prescriptive analytics are about bringing business value through machine learning.

Data — specifically, time-series data — is everywhere. It’s collected by websites, security and traffic cameras, machines, and Internet of Things (IoT) sensors, among other sources. The challenge today lies not in collecting data, but in gaining actionable insights from data with predictive analytics.



TIP

If you don’t think there is value in predictive analytics, think again. According to Grand View Research, the global predictive analytics market is growing at a 23.2 percent compound annual growth rate (CAGR) and is expected to reach \$23.9 billion by 2025.

Predictive analytics can be applied to core business processes in any industry, such as sales forecasting and demand planning in retail, electricity production and consumption in the energy sector, fraud detection and credit risk in finance, and asset health monitoring as predictive and preventive maintenance in manufacturing.

Artificial intelligence (AI) and machine learning (ML) to support predictive analytics is one of the most disruptive and innovative classes of technologies in recent years. If you don’t yet have a business strategy for predictive analytics, you’re lagging behind your competitors.

Potential business value in predictive analytics can be found in the following areas, among many others:

- » Improving customer experience
- » Reducing costs

- » Generating new revenue
- » Analyzing competitive pressures in your industry
- » Automating repetitive manual tasks

A FEW KEY TERMS TO KNOW

Here are a few important terms and concepts you should understand as you read this book:

- **Data availability:** The specific availability of observations at a certain point in time.
- **Features:** Information used by the model that is generated from the selected input variables, possibly supplemented with additional information, by doing transformations on them. These transformations include past values of certain variables, interactions between variables, moving averages of some variables, and so on.
- **Forecasting horizon:** The period of time for which forecasts are to be made.
- **Forecasting routine:** A set of forecasting situations with their corresponding data availability schemas.
- **Input variables:** All variables that are provided for an algorithm.
Selected input variables: Input variables that are used by the model and, thus, contribute to the result. They meaningfully relate to the target variable and explain part of the variance of the target variable. Synonyms include *explanatory variables*, *predictors*, and *predictor variables*.
- **Target variable:** The input variable to be modeled and forecasted or predicted to be an anomaly.
- **Time series:** A series or sequence of data points in time order, most often taken at successive equally spaced time intervals (for example, monthly, daily, hourly, and so on).
- **Time-series data:** Data following a time-series pattern or format.
- **Time-series problems:** Statistical problems or questions relating to time-series data (for example, forecast the temperature for the next few days based on historically measured temperatures; detect anomalous gas consumption based on historically measured gas consumption, temperature, and wind speed).

In the past, the field of analytics was mainly driven by data scientists and IT professionals. Augmented predictive analytics transforms and democratizes how the business user creates, uses, and analyzes predictive models for forecasting and anomaly detection. Thanks to automation, expert data scientists can now be more productive and ML model generation is easily accessible to a broader range of business users, including citizen data scientists.

Similarities can be found in the evolution from mainframe reporting to business-user-focused business intelligence (BI) tooling. In the past, IT departments developed reports for business users, causing unproductive debates about different functionalities. IT felt that the business didn't know what they were talking about, and business felt that IT overcomplicated everything and didn't deliver. This is the classic business/IT alignment problem. With augmented predictive analytics, you empower business users and citizen data scientists while enabling your data scientists to be more productive.

Understanding the Value Proposition of Data and Analytics

Getting buy-in for your data and analytics projects requires business and project stakeholders to be able to clearly, consistently, and frequently articulate the value proposition — not just the goals or strategy — of the project. The value proposition must be explicit, understood, and agreed upon by all parties.

Value propositions for data and analytics projects generally consist of one of the following:

- » **Business utility:** The results of the project are immediately accessible to stakeholders to be used as a tool for the business, as needed.
- » **Business enabler:** The results of the project inform business decisions and help business leaders make the right decisions at the right time based on the best available information.
- » **Business driver:** The results of the project uncover hidden insights and drive new opportunities for the business.

These value propositions aren't mutually exclusive and don't necessarily represent a level of analytics maturity within an organization. Different stakeholders will use data and analytics differently, depending on their individual goals and requirements.

Defining a Business Value Model

For many years, traditional measures of business value were based on established accounting and finance metrics, such as net profit and return on investment (ROI). Lagging indicators such as these are no longer effective in the age of digital transformation, in which real-time predictive insights drive greater business agility and uncover new opportunities.



TECHNICAL
STUFF

A *leading indicator* is a predictive measurement. For example, the percentage of people wearing face masks in a city could be a leading indicator for the risk of COVID-19 infection. The number of patients infected with COVID-19 in hospitals is a *lagging indicator*.

Let's explore how leading indicators can help you define your predictive analytics projects. The key is to find use cases in your business where predictive analytics can really make a difference, such as the following:

»» Reducing cost

- Eliminating waste
- Improving efficiency
- Optimizing resources

»» Improving quality

- Enhancing customer experience
- Reducing variability or unplanned events

»» Generating revenue

- Improving flexibility
- Increasing agility
- Removing dependencies and inertia
- Enabling new products and services

»» Reducing risks

- Improving stability
- Improving visibility and understanding
- Reducing complexity

Tangent Works has developed the Tangent Information Modeler (TIM) Value Framework. A simple model (with the elements listed earlier) is available, as is a more detailed model, which includes business focus, aggregates, and key performance indicators (KPIs). This framework can help you identify the value of the predictive analytics business use cases that you want to implement. It also provides a great way to communicate the value within the organization to help you land the project.



TIP

Consider the following use cases and business applications for your organization:

- »» Demand/capacity prediction
- »» Predictive asset management
- »» Fraud/anomaly detection
- »» Marketing/social analysis
- »» Pricing optimization

A MACHINE LEARNING HISTORY LESSON

Years ago, programmers used assembler language on mainframes to develop code. Fortunately, the market has been democratized, and powerful tools have become available. Now, the same trend has occurred with data analysis.

The first era of ML — call it ML 1.0 — was all about handcrafted models built in the backroom by teams of experts. It took weeks, months, or even years for data scientists and engineers to gather their training data, engineer their features, select algorithms with intimidating names (like “recurrent neural nets,” “support vector machines,” and “extreme gradient boosted trees”), tune hyperparameters, test, validate, and then iterate through it all until they finally had a model that could be deployed by another set of experts. It’s no wonder that 95 percent or more of the models never made it all the way through this process.

In the era of ML 2.0, automated machine learning (AutoML) arrived, promising to automate (at least part of) the manual processes of ML 1.0. You could now load your data into an AutoML system which, through brute force, heuristics, and a boatload of compute power, could crank through dozens of feature engineering and preprocessing steps, try tens or even hundreds of fancy algorithms, perform a grid search to find the optimal hyperparameters, and then retrain and validate your model across multiple data partitions. If your data was small enough (say, a few thousand rows) and your algorithms were fairly simple, it would take tens of minutes to get a model — enough time to sit back and enjoy a cup of coffee or two. However, if your training data was larger, your waiting time would be a few hours — enough time for far too much coffee. AutoML was a huge improvement over the weeks and months of model building in ML 1.0, but nowhere near the speed to insights that's necessary in the BI world of today. Specifically, in the world of time-series data, waiting a few hours on insights is synonymous with having outdated insights. Decisions need to be made in (near) real time and are based on these insights.

This brings us to instant machine learning (InstantML), ML 3.0, which can spark a revolution in the way that ML is used within organizations, similar to the shift seen in BI. InstantML has been introduced by Tangent Works as a completely new ML paradigm. It doesn't just automate data analysis processes — it flips the processes on their heads. In short, instead of the time- and resource-intensive multistep process of AutoML, InstantML engineers features and applies a highly efficient algorithm in a single step. This yields a model and accompanying insights in seconds. It's by far the fastest path from data to predictive value.

IN THIS CHAPTER

- » Moving predictive analytics projects from experimentation to production
- » Ensuring high-quality and available data
- » Building a data strategy
- » Changing the approach: from handcrafted modeling to software as a service
- » Handling structural changes

Chapter 2

Understanding the Challenges of Implementing Predictive Analytics

Many organizations have limited skills and experience to effectively implement machine learning (ML) and predictive analytics. As a result, ML and predictive analytics are often seen as an experimental playground. Implementations don't make it to production; instead, they get labeled as a "nice try." Organizations often view ML as an innovative and complex domain. Although many people realize that it has great potential, this potential is often deemed to be a long way off. All this contributes to ML projects getting stuck in the experimental stage.

Predictive analytics projects often face several challenges, including the following:

- » **Working with limited staff and skills:** It's estimated that only 23.7 percent of organizations have data scientists; among these data scientists, only 20.5 percent have an extensive educational background in ML.
- » **Struggling to build and set up ML operations (MLOps):** As the number of models grows, it's essential to keep track of your models and have controlled change and configuration management. (A parallel challenge exists for DevOps in the domain of software development.)
- » **Understanding the value of predictive analytics:** Turn to Chapter 1 for more on this challenge.
- » **Finding the right use case and determining the business value and priorities:** Turn to Chapter 1 for more on this challenge.
- » **Managing integration complexity:** Predictive analytics solutions need to integrate in the overall IT landscape of an organization (more on this in Chapter 7, where the importance of smooth integration in existing tools and applications is discussed).

Several other key challenges in predictive analytics projects, discussed in this chapter, include the following:

- » Data quality and availability
- » Project and operation governance
- » Lack of a data strategy
- » Being stuck in handcrafted modeling and failing to move to production
- » Structural changes in the environment, rendering existing predictive models obsolete

Data Quality and Availability

Data is the new oil in the information age but many organizations struggle to maximize the value of their data. Merely collecting data isn't enough; before you can start to analyze the data,

it needs to be organized and available in a clean format. Organizations must have a strategy that ensures data for the predictive analytics processes is collected and prepared smoothly. Many algorithms rely on high-quality and clean data, and even the best algorithm can only get out of the data what's in it. This illustrates the importance of having high-quality data available when looking to derive insights from it.

Project and Operation Governance

As mentioned in the introduction of this chapter, many predictive analytics projects never get past the experimental stage. Many projects initially look promising but fail to deliver real business value. These failed projects share a few common characteristics, which you should be aware of to avoid this trap.

Consider using the following approach for your predictive analytics projects:

» **Start early, start smart.** It pays to do some experimenting to discover what works for your organization. Many projects are too large and get stopped before they can deliver value. For example, an aircraft manufacturer might try to take on a predictive maintenance project for commercial airliners. The typical airliner has more than 25,000 sensors and 4 million parts and generates over 70 terabytes (TB) of data per flight. It's easy to see how one could get lost in the complexities of such a project, especially when there's a lack of previous experience.

Instead, the manufacturer may consider a smaller project to get started — for example, focusing on a subcomponent of the aircraft. A quick return on investment (ROI) from such a small starter project will help pay for subsequent projects. Additionally, this success story will help to change or set the mindset about predictive analytics projects, gaining traction for subsequent initiatives. **Remember:** Don't try to boil the ocean — take it one kettle at a time.

» **Clearly define your project.** We're well past proving the efficacy of ML. Identify a project that will immediately deliver business value and can be easily understood by all the stakeholders. When defining the project, design top-down and build bottom-up (see Figure 2-1).

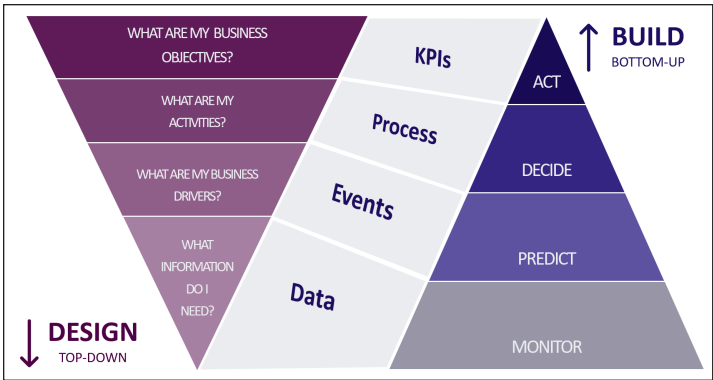


FIGURE 2-1: Let your executives and business leaders define the project; then let your data scientists execute it.

In other words, your organization’s executives and business leaders should first define the project in terms of the business benefits it will deliver. Then let the data scientists figure out how to use the available data to deliver the desired results. This approach ensures stakeholder buy-in right from the start. **Remember:** It’s a business project, not a data-science experiment. Let your business leaders do what they do best, and let your data scientists do what they do best.

- » **Track success.** Set measurable key performance indicators (KPIs) and success targets. KPIs are things you measure (such as cost savings); success targets are levels you need to achieve to be successful (such as a 20 percent total cost reduction in a process). In tracking success, also remember to consider all the benefits of the project. For example, reducing downtime in a specific piece of equipment may result in a direct saving of \$5,000 per hour because of the avoided downtime, but if an outage would also result in shutting down the entire production line at a cost of \$200,000 per hour, there is a substantially larger benefit to the project.

There are four types of KPIs — technical, qualitative, performance, and financial. Ultimately, business leaders care most about financial KPIs.



REMEMBER

Lack of a Data Strategy

For a predictive analytics project to achieve its goal(s), developing accurate models is not sufficient. Successful initiatives start with a well-thought-out problem statement, including a clear business question, and go beyond modeling and analyzing to effectively communicate the insights to stakeholders.

Data scientists mainly focus on data and technology — and they should — but there's more to a successful initiative. A successful project starts with a clearly defined use case, asking one or a few specific questions. Data scientists can then move on to look at what data is needed for answering these questions and find out if this data is available. Possibly, the questions will need to be refined before the problem statement is aligned with the available data (refer to the example in Figure 2-1). After this alignment is ensured, the data scientists can start cleaning and preparing the data, building the algorithms, and creating the models. But even after all this, the process isn't complete. The insights need to be communicated to the stakeholders in a structured and understandable way. For stakeholders to base their decisions on these insights, a sense of trust in these results is required. This trust is largely dependent on clear and effective communication from the data-science team to the business stakeholders.

Being Stuck in Handcrafted Modeling and Failing to Move to Production

Many organizations fail to see the value of predictive analytics that is available as software as a service (SaaS) and instead get stuck in developing custom solutions. Data scientists often stick to the modeling techniques they learned in school or have experience with. Many modeling techniques can be used across a range of problems, but that doesn't mean they deliver the ideal solution. Still, learning new techniques takes time, and the pressure to deliver results is often high. Frequently, settling for an acceptable result takes priority over optimizing the approach, even when developing custom solutions.

Complex projects in niche applications may require custom solutions to be made by experts, but the vast majority of projects

don't fall into this category. There are synergies to be found in using available services, especially because it's nearly impossible to have expertise in every domain. To take advantage of these synergies, many organizations need to change their approach to predictive analytics projects. By allowing data scientists to focus on a particular type of problem and come up with an optimized approach for these problems, they can build solutions that are more broadly applicable.

For example, data scientists can come up with a solution for time-series analysis in general, instead of creating a model that can be applied to one particular time-series use case, such as price forecasting for one specific product. When these solutions are offered as a service, organizations looking for them can make easy use of them. To expand on the previous example, an organization could use the service for all its time-series analysis projects. The domain experts in the organization can apply their knowledge to feed the right input data into the service and to interpret the results that are produced.

Structural Changes in the Environment Rendering Existing Predictive Models Obsolete

You're proud of your handcrafted predictive model, the hard feature selection and engineering journey, the model selection, and tuning and testing. Then a structural change happens, caused by something like a pandemic, regulatory changes, or market evolution. Suddenly, your model is no longer suitable. Back to the drawing board. To address this challenge, organizations must redefine their current ML models, strategies, and value propositions.

Let's take a closer look at one of these examples: COVID-19. ML models must now deal with new predictors, new data, and new data availability challenges, as well as the absence of reliable and relevant historical data in the wake of a pandemic that is unprecedented in recent history.



TIP

Wondering what these data availability challenges entail? Chapter 4 takes a closer look at the difficulties of time-series analysis, explaining data availability in more detail.

Structural change renders existing models obsolete and destroys the business's trust in predictive analytics. Augmented predictive analytics can offer functionality to deal with these changing conditions, allowing organizations to quickly adapt to the changing circumstances and strengthening the belief in the results or ML initiatives.



TIP

Find out more about structural changes and the challenges they introduce in Chapter 4.

IN THIS CHAPTER

- » Checking out use cases in the retail industry
- » Lighting up time-series data in the energy sector
- » Calculating the value of predictive analytics in finance
- » Increasing productivity and uptime in manufacturing
- » Communicating the value of predictive analytics in telecom

Chapter 3

Exploring Business Use Cases

Time-series data is everywhere — from banking, education, and healthcare to manufacturing, retail, transport, utilities, and many other industries. In this chapter, you explore business use cases in different industries and how the Tangent Information Modeler (TIM), discussed in Chapter 6, helps businesses maximize the value of their data with predictive analytics.

Retail and Consumer Packaged Goods

The retail and consumer packaged goods (CPG) industry is characterized by relatively low margins, high seasonality, demand variability, inventory risk, and influence of consumer sentiment. These characteristics amplify the importance of predictive analytics for key activities such as sales forecasting, demand planning, supply chain management, and inventory optimization.

To stress the importance of these activities, consider the following research about consumer perspectives at the start of the COVID-19 pandemic, conducted by BlueYonder:

- » Eighty-seven percent of consumers have experienced out-of-stock products, both in-store and online.
- » Seventy-nine percent of consumers were more likely to buy the same product from a different retailer if a desired product was out of stock.
- » Seventy-nine percent of consumers were more likely to buy a different brand of a product from the same retailer if their desired brand of that product was out of stock.

The bottom line — which affects *your* bottom line — is that inventory availability supersedes brand loyalty. Yet inventory availability is dependent upon numerous variable factors across the entire value chain (see Figure 3-1). Reducing and accounting for this variability through better planning based on better forecasts (from better forecasting models) is the key to optimizing inventory, ensuring the right product is available at the right price at the right time, thus meeting customers' needs.




FIGURE 3-1: Each link in the value chain must serve a demand in the quickest possible way.

Sales forecasting

Forecasting sales of a product or service plays an important role in the life cycle of almost every retail company. The estimation of future sales can drive plenty of management decisions, such as efficient inventory management, prevention or early detection of potential issues, price setting, and marketing.

Accurate sales forecasts in the retail industry can be the difference between profitability and insolvency for retailers in today's hypercompetitive markets. Sales planners frequently rely on forecast-driven tools to adjust levers such as product pricing and the timing of promotions. This process is complex, because these techniques need to be applied dynamically across different levels of product hierarchy, geography, and other dimensions. However, many companies are still using relatively rudimentary forecasting techniques, which can adversely affect the accuracy of the resulting forecasts. Moreover, many enterprise-facing tools are designed with inefficient workflows, reducing the ability to do effective analysis for end users such as financial planning and analysis (FP&A) professionals and sales-planning teams.

Typical data that can be used as input in sales planning include the following:

- » Historical sales data (usually segmented by product hierarchy, geography, and so on)
- » Regional store information
- » Local demographics
- » Level of competition
- » Indicators of consumer demand, industry performance, and economic performance
- » Pricing and promotion start and end dates



TIP

Machine learning (ML) techniques provide the most up-to-date approach for accurate forecasting, but they can be time consuming to implement. TIM's instant ML (InstantML) and real-time instant ML (RTInstantML) capabilities allow analysts to easily apply forecasting models to any time-series data and quickly iterate on planning scenarios.

Demand planning

Another common task for retailers is to forecast future demand. For example, without a qualitative forecast, it can be very difficult to assume the right amount of stock that should be available. Retail forecast accuracy is negatively affected by rapidly changing conditions and forecast errors reach 30 percent or more, on average. Demand planning is a challenging use case, but one with lots of potential for improvement.

Typical data includes past sales volumes, supplemented with data regarding commercial activities and external factors impacting sales volumes. Data sets often vary depending on product, sector, or even geographic location, resulting in a cumbersome and complex model-building process.



TIP

TIM provides automated selection of the right input variables with an explanation of the impact of each *predictor* (selected input variable), allowing for further refinement or data sourcing. Automated model tuning based on internal and external changes enables greater responsiveness and results in less waste due to inventory scrapping (especially for perishable goods), as well as fewer lost sales due to inventory shortages.

Energy Sector

From electricity generation to storage, transport, distribution, and consumption, the energy industry value chain is extremely volatile and is being rapidly transformed by global market trends such as deregulation, renewable energy, carbon footprint reduction, energy exchanges, and smart meter/grid technologies.

It's critical for many companies to accurately forecast electricity load. Electricity load is a fundamental input to operations of transmission system operators (TSOs) and is important for industrial producers to balance their decisions on electricity procurement. Owners of photovoltaic (PV) plants, electricity traders, and system regulators need accurate forecasts of production of the PV plants, for different time horizons and different granularities, to optimize their maintenance, trading, and regulation strategies; the same goes for wind production. These examples only scratch the surface of the widespread presence of time-series data in utilities.



REMEMBER

Modeling time-series data for the energy industry supports key decision-making that can affect short-term supply-and-demand planning, energy efficiency, spot market futures, energy production, and long-term capacity management.

Energy consumption

Industries, companies, cities, and households all consume energy. Whether opting for electricity, gas, or thermal power — or, more likely, a combination of them — the need for energy is all around us. Both consumers and producers can benefit greatly from accurate estimates of future consumption, not in the least because extreme volatility of wholesale prices forces market parties to hedge against volume risk and price risk.

The ability to accurately forecast future energy consumption is a key determining factor in the financial performance of energy industry players. Business decisions based on incorrect energy volume estimates can be costly.

Consider the following example: Looking at a rough estimate of savings from a 1 percent reduction in the mean absolute percentage error (MAPE) of the load forecast for 1 gigawatt of peak load can save a market player approximately:

- » \$600,000 per year from long-term load forecasting
- » \$350,000 per year from short-term load forecasting
- » \$700,000 per year from short-term load and price forecasting

The value of ML is clear, but it must be weighed against the cost and effort it requires. To achieve accurate forecasts, relevant predictors must be used. Explanatory variables in energy consumption use cases include the following:

- » Historical load data (in different levels of aggregation)
- » Real-time measurements
- » Weather-related data (such as wind speed)
- » Calendar information
- » Day/night usage patterns



TIP

TIM automates model generation of accurate forecasting models and tells you which input variables have the highest relevance in calculating the forecasts. With its InstantML approach, TIM creates these models in mere seconds rather than the days or weeks common for handcrafted models and AutoML approaches. The

scalability of TIM's model-generation process allows for hundreds of models to be generated at the same time.

Wind production

Although ecologically friendly and quite popular, wind production is a volatile source of energy. Wind is difficult to predict, and generated electricity is hard to store. Wind production use cases rarely center around a single windmill or even a single wind farm; instead, they often involve a large portfolio of wind assets that are strategically located to take advantage of favorable climate conditions. The larger the portfolio, the more difficult it is to manage and obtain an optimal dispatch and exposure to the electricity market.

Besides great opportunities for balancing the grid and forecasting production, predictive analytics use cases for wind production also involve a lot of predictive maintenance. Explanatory variables can include the following:

- » Weather-related data (particularly wind speed, temperature and humidity)
- » Technical information (such as wind turbine types and height)



TIP

TIM can automate the modeling of complex wind (and solar) production scenarios. Moreover, TIM allows for blended forecasts that unify high-quality intraday modeling and day(s)-ahead modeling. TIM's output consists of the forecasted wind production in the same unit of measurement (typically, kilowatts per hour [kWh] or megawatts per hour [MWh]) and granularity as the input data, over the desired forecasting horizon.

Solar production

Many different parties are impacted by the production of PV plants, from owners to electricity traders to system regulators. This production has an impact on multiple domains, such as maintenance, trading, and regulation strategies. However, the high short-term volatility in solar production makes balancing the grid a difficult task. Moreover, a single impacted party often has interests

in a large portfolio of solar assets, which may consist of different sizes of plants at different locations. Inaccurate forecasts can result in significant financial penalties, whereas improvements in forecasting accuracy can lead to significant financial gains. Large portfolios with important impacts require consistent and scalable forecasts. Achieving high accuracy isn't the only challenge, though — data availability for these large portfolios of volatile assets can be a problem, too (see Chapter 4).

Several different variables can be explanatory in this use case and should be included as inputs into the model-building scenarios, when possible. These variables include the following:

- » Weather-related data such as the global horizontal irradiation (GHI) and the global tilted irradiation (GTI)
- » Position of the sun
- » Location of the PV plant(s)
- » Direct normal irradiance (DNI)



TIP

TIM can handle different data availability situations either by allowing the user to account for the situation in the relevant model-building definition (InstantML) or by building and deploying models ad hoc (RTInstantML), taking into account the current data availability situation. (I discuss both InstantML and RTInstantML in Chapter 5.)

Finance

In finance, entire companies are in business based on their ability to foresee future stock prices. The finance industry has many ML use cases, from pattern recognition in hedge fund management to supporting quantitative strategies for trading in capital markets. Plus, many financial institutions use their time-series data when deciding future interest rates, when attempting to prevent fraud, when deciding which loan requests to approve, and in many other scenarios.

Mortgage prepayment and credit

The mortgage industry generates vast quantities of data that is highly relevant to profitability and risk analysis. Some of this data is in the possession of users, some of it is publicly available, and some of it can be acquired. However, the velocity of change in the mortgage industry is outpacing the abilities and reach of existing prepayment and profitability models. With more than \$16 trillion total outstanding mortgages, the U.S. mortgage market requires a scalable and accurate forecasting and modeling solution.



TIP

TIM enables users to iteratively forecast mortgage prepayment, delinquency, and default so that investors, servicers, lenders, and other stakeholders can evaluate and quantify the valuation and credit risk of their mortgage assets.

Manufacturing and Supply Chain Management

Industry 4.0 trends including cloud platforms, the Internet of Things (IoT), big data, edge computing, and the proliferation of mobile devices are driving a growing need for ML and predictive analytics in manufacturing.

PwC predicts that ML adoption, specifically to improve predictive maintenance in manufacturing and supply chain management, will grow by 38 percent. Similarly, McKinsey Consulting forecasts that ML will reduce supply chain forecasting errors by 50 percent. In a recent survey by McKinsey Consulting, supply chain management and manufacturing were identified as leading industry use cases for cost reduction through artificial intelligence (AI) applications.

New opportunities for AI and ML applications in manufacturing and supply chain management include demand, supply chain, production, maintenance, and life-cycle planning, as well as quality monitoring and predictive analytics. The benefits include reducing cost, customer churn, and waste while increasing agility, efficiency, productivity, and safety.

In the following sections, I take a look at a few specific use cases.

Predictive maintenance and asset health monitoring

Getting the most out of your production assets, especially when working within operational constraints, is key to increasing productivity. As the manufacturing industry has become more digitized, time-stamped data has proliferated. With billions of connected sensors being deployed in smart manufacturing plants, this trend will continue to accelerate. Many core decision-making processes that used to be based on relatively static information are increasingly based on more dynamic, real-time streaming data.

Highlighting anomalies in time-series data from multi-variate sensor readings helps to alert operators about potential equipment failures during production runs. These signals can then be analyzed and used as indicators for potential performance degradation or equipment failure.

Input data for predictive maintenance and asset health monitoring typically consists of raw time-series data from programmable logic controllers (PLCs), supervisory control and data acquisition (SCADA), sensors, maintenance scheduling systems, and condition monitoring processes. Some data examples include the following:

- »» Vibration
- »» Temperature
- »» Revolutions
- »» Pressure
- »» Quality
- »» Past error codes
- »» Future condition monitoring alerts
- »» Past and future maintenance schedules
- »» Past maintenance information
- »» Past and future equipment operations schedules



TIP

TIM's anomaly detection capabilities increase your return on assets by reducing unplanned maintenance and increasing equipment uptime. Users can create and deploy models at all levels of manufacturing control operations including field (sensors), direct control, plant supervisory, production control, and production scheduling.

Asset failure analysis

Complex and distributed equipment, such as differently configured pumps or compressors installed across the globe, fail for many reasons. Some failures are due to normal wear and tear; others may be due to local operating conditions or the specific configuration of an asset.

Gathering data through industrial IoT (IIoT) platforms and performing anomaly detection enables manufacturers to predict imminent failures and perform root cause analysis when anomaly detection leads to explainable forecasts. This information leads to faster resolution of issues and allows research and development (R&D) to analyze failures to improve reliability. This is particularly important when service contracts require manufacturers to bear at least some of the cost for maintaining equipment and guaranteeing uptime.

Typical time-series data sets used in asset failure analysis consist of computerized maintenance management system (CMMS) data combined with IIoT data, as well as external elements such as operating conditions (for example, weather, vibrations, speed, and so on).



TIP

TIM's forecasting and anomaly detection capabilities produce accurate and explainable results that can be analyzed by technical maintenance teams to quickly pinpoint root causes, as well as R&D teams to determine structural improvements to the equipment.

Supply chain strategy planning

Modern supply chains are highly complex and interdependent, requiring agility and velocity to keep businesses on track. Traditional AI and AutoML approaches are too expensive, slow, and difficult to adapt. Combining business data and market prognosis scenarios with RTInstantML enables organizations to create new, improve existing, and evaluate more advanced what-if scenarios

and simulations for supply chain strategic planning and business transformation.

Some examples of supply chain strategy planning processes that can benefit from InstantML forecasting are strategic budgeting exercises, business transformation and design initiatives, strategic product life-cycle planning and optimization, and product and product maintenance design processes. Typical inputs include the following:

- » Historical demand, supply, production, prices, costs, and strategic performance data
- » External data, such as weather, sales periods, global and regional disruptive events, sales campaigns, and product introduction information



TIP

TIM can be used for complementing what-if and simulation scenarios for budget exercises, adapting maintenance for product support strategies, running forecasting and anomaly detection on digital twins in product design, doing risk assessments in your business transformation process, and more.

Telecom

Telecommunication systems are coursing with time-series data. For example, a mobile operator with millions of subscribers will generate tens of millions of call records daily and billions of rows of call data over the course of a year. In turn, each of these calls is routed through a myriad of network nodes, switches, routers, servers, and other equipment, each of which generates huge volumes of telemetry data that are organized by time. Telcos now carry about 200 times more daily Internet data traffic than voice traffic (calls), and the total traffic that telcos are dealing with has increased over tenfold in the last five years alone. You can truly say that telecom companies have a big data problem — and opportunity!

The time-series data that telcos have at their disposal is a veritable goldmine of riches once ML is applied. Hidden in this treasure

trove are insights that can help telcos enhance customer experience, improve their network quality, and capture and retain more subscribers.

Testing customer loyalty with experience issues

As a telco operator, your subscribers expect your incredibly complex and expensive network to just work. Of course, most of the time, it works great — but subtle, hard-to-detect anomalies can impact your service quality and negatively affect the customer experience of your subscribers. These issues test your customers' patience and could potentially lead to subscriber churn; as a result, loyalty refunds, credits, and incentives to keep them happy eat into your profit margins.

Time-series network analytics provides the ability for network operations and engineering teams to predict trends and detect anomalies in the network performance and proactively improve the service quality of their networks. Using predictive analytics on the network telemetry — including network performance, Internet usage, clickstreams, and traffic-flow data — enables operators to identify underperforming components of the network and improve the quality of experience for subscribers before they may even realize there is an issue.

5G: The next telco battleground

According to the Global System for Mobile Communications Association (GSMA), a telecom industry association representing more than 1,200 companies, by 2025, 40 percent of the global population will be covered by 5G networks. 5G presents an incredible opportunity for mobile operators to capture market share, extend their offerings with innovative new products and services, connect more of the world's population than ever before, and enable extraordinary technological and societal advancements like autonomous vehicles.

This means, of course, that individual operators need to make careful decisions about when, where, and how to invest in 5G, as well as how to monetize the capabilities that they're investing in.

One of the most exciting aspects of 5G for individual operators is very practical. It enables something referred to as *network slicing*, which allows for telcos to create virtual slices of the network to serve specific customers, application types, or other dimensions. These slices can be tailored to very specific demands on the network and service-level requirements versus the “one-size-fits-all” approach of the past. Imagine different slices of the network dedicated to autonomous vehicles, smart city infrastructure, emergency response, streaming gaming users, and more. As an operator, you can dynamically create and optimize your 5G services to serve these slices individually and develop specific pricing structures for these individual markets accordingly.

The question this raises is, how can operators efficiently create, support, and price these network slices? This is where ML comes in. The behavior of each slice can be analyzed, and resource demands can be forecasted. This provides the operator the ability to adjust the capacity of the slice dynamically to respond to forecasted events and correlations between events (perhaps with other slices of the network). Plus, these forecasts can help operators efficiently price their 5G services to balance the cost and value that each slice is delivering. For telcos struggling with growth and profit optimization, this is a big deal.

- » Accounting for changing circumstances
- » Addressing structural changes
- » Dealing with data availability
- » Building multipoint forecasts

Chapter 4

Why Time-Series Forecasting Is Difficult

In this chapter, you learn about some of the unique challenges of time-series forecasting, such as changing circumstances, structural changes in the data, the availability of data, and creating multiple forecasts for multiple forecasting horizons in multiple situations.

Circumstances Change Over Time

Events represented by time-series models are dynamic in nature. A model that worked yesterday may no longer be valid today. For example, in finance, an investment model may be based on a portfolio of assets that changes rapidly as market conditions evolve. In manufacturing and energy, production models may change quickly due to a change in orders affecting demand or unexpected downtime affecting capacity.

When training a machine learning (ML) algorithm to play a game of “Go” or how to distinguish between images of lungs to detect lung cancer or detect or recognize fake news articles among thousands of new articles created every day, more data samples typically enhance the performance of your model. However, when

modeling time-series data, this is generally not the case. New observations in time-series data may make your model useless from one hour to the next because the underlying process can suddenly change. This forces people to repeat the model-building process, though this is clearly not an optimal solution, because it can take up much of their valuable time. It also requires people to critically evaluate which data set, over which time period, should be used as training data for the model-building efforts. Additionally, these new situations often require the identification of new significant features rather than a different modeling technique or slightly adjusted hyperparameters. So, people need to repeat even more of the process because they're forced to return to the feature recognition stage before going on to retrain their models. Understanding when you should rebuild your models along with what data you should use when rebuilding is crucial.



TECHNICAL
STUFF

A model *hyperparameter* is a characteristic or element that is external to the model; its value is used to control the ML process itself. A model *parameter* is a characteristic or element that is internal to the model; its value can be estimated from data.



REMEMBER

The Tangent Information Modeler (TIM), discussed in Chapter 6, empowers users to easily adapt to new situations by allowing models to be continuously recalibrated or rebuilt. Whereas model recalibration only adjusts the model's parameters and leaves the model's structure (features) intact, model rebuilding starts by identifying new features and then builds a completely new model.

Structural Changes

There are many different types of structural changes in time-series models, such as changes in mean, variance, and seasonality, which can fundamentally alter the way in which a business model or economy functions. Structural changes can sometimes be so significant that the data after the change appears to be a completely new and different time series compared to the data before the change.

Figures 4-1 and 4-2 show two examples of dynamic time series containing structural changes. Figure 4-1 shows an example of gas consumption for heating with a clear seasonal pattern. Figure 4-2 shows an example of the impact of the COVID-19 pandemic on travel volume. This structural change is so fundamental that it effectively seems to split the time series into two different time series.

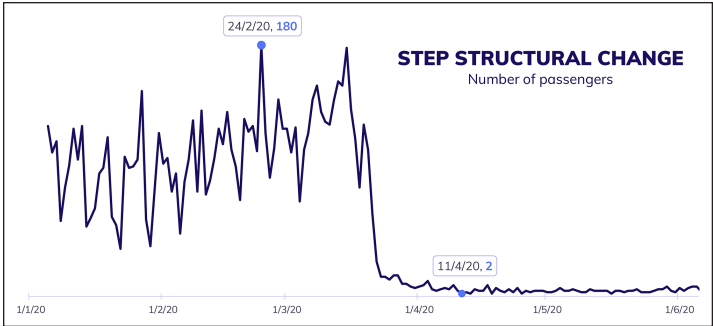


FIGURE 4-1: A seasonal structural change in gas consumption for heating.

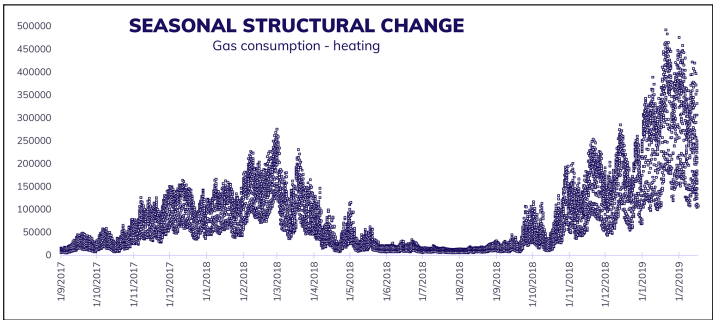


FIGURE 4-2: A structural change in the number of passengers.



REMEMBER

For time-series problems, setting up a continuous model-rebuilding pipeline is often necessary to address changing circumstances and to maintain the required level of accuracy. Swift and easy model-rebuilding capabilities, such as TIM’s real-time instant ML (RTInstantML) capabilities, greatly contribute to this process. On top of this, such capabilities empower users to handle structural changes in their data.

Data Availability

Another challenge in time-series forecasting is the availability of data for your model, which may vary at times. Certain data that was used for building the model may not be available when a forecast needs to be made, so the model can't be used for forecasting. For example, in Figure 4-3, the following relationship is identified using historical data from January 1, 2017, to January 9, 2017:

$$\text{Electricity Consumption (t)} = 10,000 \times \text{Average Temperature (t)} - 5,000 \times \text{Climate Protests (t)}$$

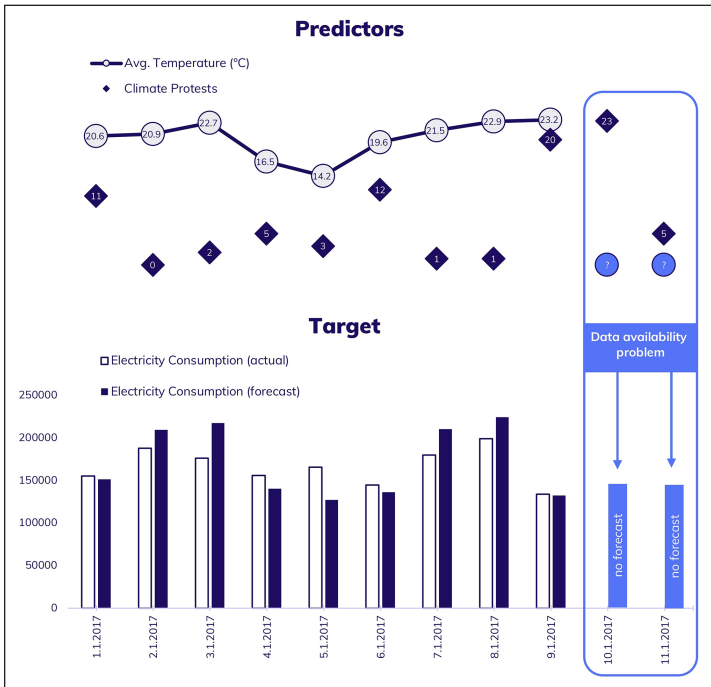


FIGURE 4-3: A data availability problem. A linear model needs data points that were available for training to be available for application (for example, when forecasting for January 10 and 11).

On January 10, 2017, the values of one of the explanatory variables (Average Temperature) are unavailable, so it isn't possible to use the identified model to calculate a forecast (see Figure 4-3). Repeating an automated ML (AutoML) search with the new constraints for

calculating Electricity Consumption (t) (namely, including Average Temperature [t] available up to one day ago and Climate Protests up to the current day) is a time-consuming option with significant computational cost. Doing this each time the data availability changes is nearly impossible.

Because TIM's RTInstantML allows for building a completely new model each time a forecast is required, there is no need to set up a particular data availability scenario. TIM looks at the data availability as it is in the data set at the time of forecasting and then builds a model based on this data and applies the model in a single pass through the data, delivering a forecast using all relevant available data in just seconds to minutes.

This approach solves the problem of (unexpectedly) unavailable data. It also seizes the opportunity that's provided in case more data is available than expected. Although this second scenario (more data available than expected) would not cause a problem for models built with AutoML approaches, the models wouldn't be able to make use of this advantage without retraining.

Need for Multiple Forecasts over Multiple Time Spans

Multipoint forecasts are those where prediction horizons include more points to predict. Multipoint forecasts are required in many industrial verticals. For example, you may need to forecast the price of a commodity at 9 a.m. for each hour that day from 10 a.m. to 9 a.m. the following day, resulting in a total of 24 points. Multipoint forecasts have traditionally been addressed by multiple output models or recurrent strategies.

Building and optimizing a multiple-output model (see Figure 4-4) is intuitively harder than doing so for a single output model, because model parameters are optimized for each of the outputs, and this optimization happens simultaneously. Optimizing for several outputs simultaneously may result in a contradictory optimization problem. Multi-output models are often more complex (for example, they may have more hidden layers in a neural net or contain a larger decision tree) and interpreting information from such a model is difficult.




FIGURE 4-4: Using a multiple-output model to forecast over a 24-point horizon.

For example, it's intuitively clear that traffic patterns during the day are far more complex than traffic patterns at night. Forecasting the traffic at 3 a.m. by looking at the traffic at 3 a.m. the previous night, would probably result in a reasonable forecast. However, forecasting the traffic at 5 p.m. is more complex than just copying the traffic from 24 hours earlier: It needs to take into account whether it's a weekday or weekend, whether it's a holiday, possibly even weather-related information (because fewer people may travel when it's raining or cold). Forecasting both scenarios using the same model provides a challenge.

On the other hand, recurrent strategies (see Figure 4-5) optimize a single-output model that is then propagated in time in a recurrent manner. For example, the forecast for $y(t + 1)$ is used again for calculating $y(t + 2)$. However, recurrent strategies are prone to diverge quickly after only a few update steps. This renders them impractical for broader adoption.




FIGURE 4-5: Using a recurrent strategy to forecast over a 24-point horizon.

How does TIM approach the need for multipoint forecasts? As shown in Figures 4-4 and 4-5, both traditional approaches bring notable disadvantages with them. To overcome these disadvantages, TIM takes a completely different approach. For each point in the forecasting horizon, TIM creates a separate model (see

Figure 4-6). In doing so, no contradictory optimizations need to be done because each model needs to be optimized for just one output. Because no outputs are used in creating other outputs, the risk for propagating errors (and, thus, quickly diverging results) is also removed. All models TIM creates in this scenario are assembled into a so-called modelZOO. TIM then automatically dispatches the correct model from the modelZOO each time a forecast is desired. This approach requires the creation of multiple (sometimes many) models, so it's only possible with TIM's fast and scalable model-generation capabilities.




FIGURE 4-6: TIM's approach to forecasting over a 24-point horizon.

In many time-series applications, it's necessary to provide a multipoint forecast several times per day (for example, every hour). Each forecast typically has a unique data availability scheme. For example, at 8 a.m. the target could be available until six hours ago, but at 9 a.m. the database gets updated with actual data up to two hours ago. Models for 8 a.m. can take advantage of lagged features from $t - 6$ hours and further back in time, but models for 9 a.m. can take advantage of more recent data from the morning update (that is, from $t - 2$ and further in time). Therefore, forecasting for 10 a.m. at 8 a.m. and at 9 a.m. is significantly different.



TECHNICAL
STUFF

A set of forecasting situations with their corresponding data availability schemas is referred to as a *forecasting routine*.

For each of the challenges of time-series forecasting discussed in this chapter, TIM's approach to overcoming them is also discussed. This may have aroused your curiosity about what TIM is exactly. If so, you're in luck, because the remaining chapters of this book discuss TIM in more detail.

IN THIS CHAPTER

- » Putting a predictive analytics solution together
- » Starting from scratch with handcrafted models
- » Automating basic model-building steps
- » Taking model building to the next level with instant machine learning

Chapter 5

Exploring Time-Series Predictive Analytics

In this chapter, you learn model-building techniques to deliver a predictive analytics solution, as well as the different model-building techniques that are used to deliver a solution.

How to Build a Predictive Analytics Solution

A predictive analytics solution is built on models. Data scientists typically build machine learning (ML) models to support predictive analytics applications. As the value of predictive analytics is increasingly sought after, the number of open positions for data scientists has also exploded. In today's job market, the demand for data scientists surpasses the supply. In predictive analytics projects, business value lies in the insights that can be derived from data. Data scientists can contribute the most to this value creation when they're allowed to focus on retrieving and communicating these insights. Any tasks in their workload that can be automated — from data preparation to feature engineering and even model generation — to free up time for these high-value

tasks should be automated. In this way, organizations can ensure that the data scientists who work for them focus on gaining insights and understanding and creating business value, rather than on repetitive and tedious tasks that don't bring direct value.

MACHINE LEARNING 101

Here's a quick primer — or crash course (depending on your background) — in ML.

ML is the scientific study of algorithms and statistical models that computer systems use to effectively perform a specific task without using explicit instructions, relying on patterns and inference instead. ML algorithms build a mathematical model of sample data, known as *training data*, to make predictions or decisions without being explicitly programmed to perform the task.

ML algorithms are used in a wide variety of applications where developing an algorithm of specific instructions for performing a task isn't feasible. It's closely related to computational statistics, which focuses on making predictions using computers to analyze large sample sizes and use intrinsically intensive statistical methods.

There are two main types of ML tasks relevant to predictive analytics: supervised and unsupervised learning.

In *supervised learning*, the algorithm builds a mathematical model from a set of data that contains both the input variables and the desired output. For example, when trying to determine whether an image contains a certain object, the training data for a supervised learning algorithm would include images with and without that object (the input), and each image would have a label (the desired output) designating whether it contained the object. In special cases, the input may be only partially available or restricted to special feedback. This category falls into two subcategories:

- **Classification:** Classification algorithms are used when the outputs are restricted to a limited set of categorical values. An example of a classification algorithm is one that filters emails, in which case the input would be an incoming email and the output would be the folder in which to file the email.

- **Regression:** Regression algorithms are characterized by their continuous outputs, meaning they may have any value, possibly within a range. Examples of a continuous value are the temperature, length, or price of an object.

In *unsupervised learning*, the algorithm builds a mathematical model from a set of unlabeled input data, meaning there is no knowledge about a desired target variable. These algorithms are used to find structure in the data (for example, by grouping or clustering data points). In this way, unsupervised learning can discover patterns in the data and can group the inputs into categories.

The study of mathematical optimization delivers theory, methods, and application domains to the field of ML. An important subfield of ML is data mining, which focuses on exploratory analysis through unsupervised learning. In its application across business problems, ML is also referred to as predictive analytics.

Handcrafted Models

The first step in creating a predictive analytics solution for forecasting or anomaly detection is to prepare the data (see Figure 5-1). This includes gathering historical time-series data and cleansing it. Other steps in the data preparation process often include data normalization or standardization.

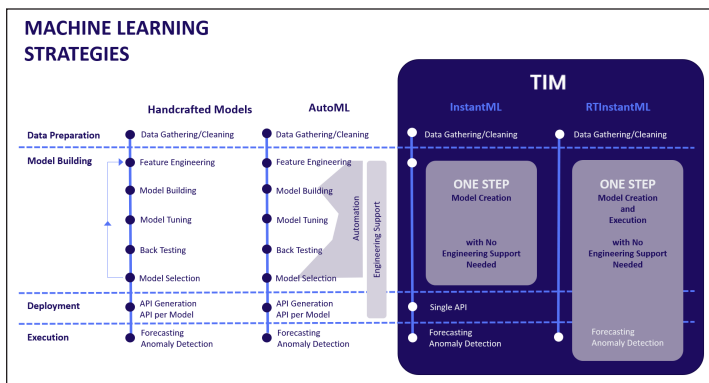


FIGURE 5-1: Current ML strategies include handcrafted modeling and automated ML (AutoML).



Normalization is a technique often applied as part of data preparation for ML. The goal of normalization is to rescale the values of numeric variables in the data set to a common range.

The next step is to build the model. Many companies have traditionally employed data scientists and engineers to build handcrafted models. Model building is an iterative process that involves feature engineering, model building, model tuning, back testing, and model selection. This process can take days or weeks and often must be repeated in order for any changes to the context in which the model will be applied.

After the model has been successfully built, an application programming interface (API) can be created for each model for deployment. Finally, the model can be applied.

Disadvantages of this approach include that it:

- »» Does not support business agility and real-time decision-making in changing circumstances
- »» Is not easily (or economically) scalable or adaptable
- »» Is expertise intensive
- »» Is engineering driven and time consuming (taking days or weeks to build and deploy)



Data scientists often spend as much as 80 percent of their time in the data preparation step. However, business value is obtained by interpreting models and results to gain insights. Freeing up data scientists' valuable time so they can move on to the tasks where business value is created — such as getting insights and gaining a better understanding of the data, models, and results — is beneficial to the business.

AutoML

Automated ML (AutoML) is the next evolution in model building. AutoML automates part of the model-building process from feature engineering to model selection (refer to Figure 5-1).

AutoML techniques train many different models and then select the most successful model. This is typically a laborious, time-consuming task that requires scanning many different

ML libraries, creating models, and tuning their corresponding hyperparameters.

Although AutoML is a step toward automation, tedious and manual feature engineering is often required, which calls for the input of a domain expert. The result is a compute-intensive trial-and-error process that can still take up hours or days of valuable time. Even so, this is a significant difference to the days or weeks of time needed for handcrafted modeling.



REMEMBER

AutoML still requires data scientists for data preparation and for engineering support throughout model building and deployment, but it does provide some scalability and accelerates the model-building process.

A New Paradigm: InstantML

In time-series modeling, identification of significant features and an overall modeling framework (how to address changing dynamics in time-series data, dynamic data availability, multipoint and multi-situational forecasts, and so on) are far more important than choosing a specific modeling technique and its associated hyperparameters.

Tangent Works has developed a new model-building framework that addresses time-series modeling challenges. The Tangent Information Modeler (TIM), discussed in Chapter 6, is an automatic model-building engine for time-series forecasting and anomaly detection. With a single pass through the data, TIM generates a high-quality model (see Figure 5-1).

InstantML is a new paradigm in model building that refers to modeling strategies that focus on identifying features rather than a modeling technique and its hyperparameters. InstantML is significantly faster than AutoML, requiring only seconds or a few minutes for building forecasting and anomaly detection models. This enables instant forecasting, where new models can be created on demand.

With InstantML, you set up your data availability scheme and TIM creates a model based on it. This process takes only seconds to minutes. You can then apply this model as many times as you want, even on new data, given that the data availability doesn't

change. Should you need a new model, because of a change in data availability, a change in the underlying circumstances, or even just because you want to, you can easily generate a new one with TIM by repeating the process.

Real-time instant ML (RTInstantML) goes one step further by unifying the steps of model building and model application (forecasting). With RTInstantML, you don't need to set up any data availability scheme as it's automatically detected from the data set that's provided. The model is trained and applied, and then discarded. With RTInstantML, your forecast is directly returned to you. The next time a forecast is needed, you can simply repeat this process.

Both InstantML and RTInstantML are explainable, allowing users to look into the models that were built. This way, users can understand the forecasts and anomaly detections produced by the models and explain them to decision-makers and key stakeholders in the business.

RTInstantML is especially advantageous if your data availability scheme changes often or if ad hoc forecasts are required. RTInstantML automates everything from model building to model application, taking input data and directly providing the user with the relevant forecast. This eliminates the need to set up a particular data availability situation, as the situation can be extracted directly from the input data set at the time of forecasting. RTInstantML doesn't require any model management or model storage, which simplifies implementation complexity and shortens time to value.

The benefits of InstantML and RTInstantML include the following:

- » **Speed:** Models and forecasts can be built on demand in seconds to minutes rather than days or weeks.
- » **Scalability:** They're highly scalable and adaptable to changes in time-series data sets.
- » **Automation:** Models are created automatically, reducing the engineering time required by data scientists.
- » **Explainability:** Human-readable models deliver truly explainable artificial intelligence (AI).



REMEMBER

Figure 5-2 compares the different ML techniques discussed in this chapter.

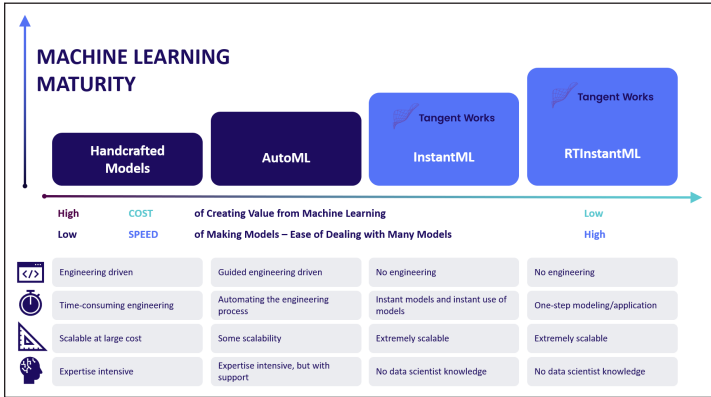


FIGURE 5-2: ML maturity.



TIP

To go beyond experimenting, businesses need to automate the model-building and deployment stages of modeling with repeatable and scalable ML.

IN THIS CHAPTER

- » Looking under the hood: the Tangent Information Modeler Engine
- » Leveraging the Tangent Information Modeler application programming interface
- » Getting acquainted with TIM Studio

Chapter 6

Implementing the Tangent Information Modeler and Instant Machine Learning

The Tangent Information Modeler (TIM) is the automatic model-building solution from Tangent Works. TIM is designed specifically for time-series data. Based on this data, TIM extracts relevant features and builds explainable forecasting and anomaly-detection models.

In this chapter, you learn about the TIM solution architecture: the TIM Engine, the TIM application programming interface (API), and TIM Studio (TIM's web-based user interface).

The Tangent Information Modeler Engine

The TIM Engine is built on a field of mathematics called *information geometry*, an interdisciplinary field that uses differential geometry techniques to study probability theory and statistics. The TIM Engine contains TIM's machine learning (ML) functionality.

This includes time-series model generation, as well as the calculation of the resulting forecasts and anomaly detections. The TIM Engine is fully containerized and can be used as a cloud service, in your own cloud, on premises, or on an Internet of Things (IoT) device on an edge network (see Figure 6-1).

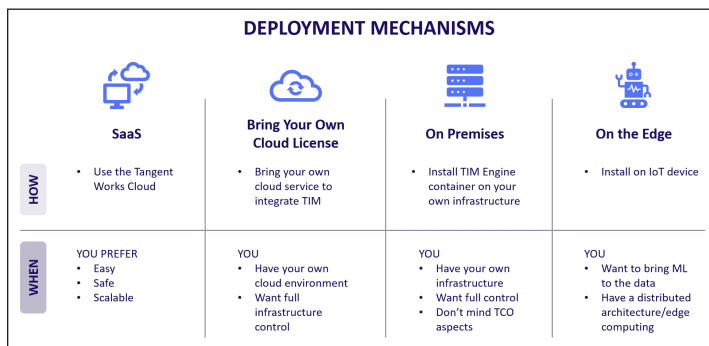


FIGURE 6-1: The TIM Engine can be deployed in various ways to meet your business requirements.

The Tangent Information Modeler Application Programming Interface

All of TIM's functionalities can be accessed via a single open representational state transfer (REST) API. This API is kept up and running at all times, automatically scaling with the number of incoming requests. Because all functionalities (model building, forecasting, anomaly detection, and so on) are accessible via a single API, TIM is very easy to use.



TECHNICAL
STUFF

Later, you'll come across different user interfaces through which TIM's capabilities can be consumed, such as TIM Studio (later in this chapter) and several platforms (Chapter 7). Under the hood, all these interfaces are just a level of abstraction added between you (the user) and TIM's API.

The high-level architecture of the TIM web service is shown in Figure 6-2. In the background, TIM constantly runs a few Workers that can facilitate all possible requests. Scalability is a very important feature of this architecture, achieved by distributing the highest computational complexity to the TIM worker units.

The number of active workers automatically scales with the number of incoming requests; when computational demand increases, additional TIM workers can be automatically started.

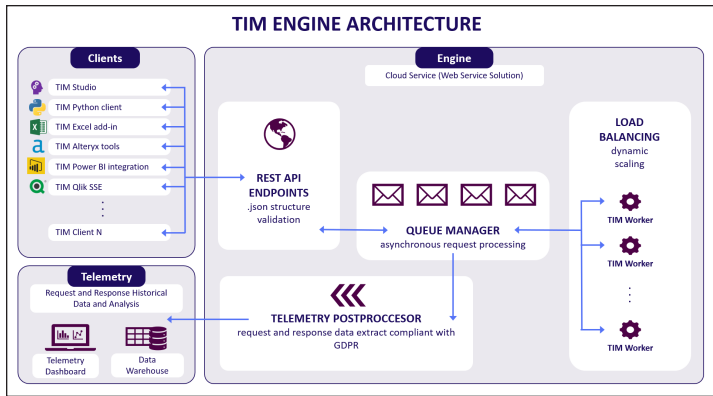


FIGURE 6-2: The TIM web service architecture.

TIM Studio

TIM Studio is a web application that offers an intuitive user interface to the TIM Engine for forecasting and anomaly detection. It's designed to serve as an ML development and operations (MLOps) platform for time-series modeling. TIM Studio is a productivity tool that brings TIM's predictive and prescriptive analytics capabilities to nontechnical users, citizen data scientists, and data scientists to create models, test results, understand their data, and apply models in production.

TIM Studio consists of the following main parts:

- » **Data sets:** Data sets hold data that are required for both model building and forecasting. TIM Studio allows you to explore data sets in more detail. Besides plotting data on charts and displaying them in tabular form, it also provides key statistical insights. If your data contains missing values, you can quickly find out where they occur.
- » **Use cases:** Use cases collect all work (including experiments) related to a certain business objective in one place. You aim for your efforts to eventually deliver benefits to business processes, so a use case can be described with a business objective, value, key performance indicator (KPI), and so on.

- » **Experiments:** Experiments belong to a use case and are tightly coupled with a concrete data set. In addition, they contain specific parameters for the TIM Engine. TIM Studio helps you keep track of adjustments that were made to various parameters and evaluate and compare results. In doing so, TIM Studio supports an iterative approach to experimenting.
- » **Production setups:** Production setups are based on a certain iteration of an experiment, namely the one that delivered the best results during back testing. TIM Studio offers a simple way of transitioning from experimenting to a regular forecasting process.
- » **Forecasts:** Forecasts are values calculated with the use of a certain production setup. Every forecast is stored in TIM Studio, so you can get back to it at any time.

Let's talk about the screen that you'd likely spend the most time with: the Experiment Workbench. The Workbench is a screen that allows you to get insights into the models generated by TIM and to evaluate their quality. This is done in a process called *back testing*. Back testing involves building a model on historical data and then assessing its quality by applying it to "unseen" data. This gives you an impression of how the model would behave in production (that is, with new data).

Users can alter the process of model generation by adjusting various settings, which are accessible via the Settings task pane, where they're grouped by category. For example, besides setting the length of the forecasting interval, it's also possible to set model-building and testing intervals or to adjust which mathematical transformations are used in the model-building process. Most of the settings can be left in the automatic mode. It's up to users to decide if they want to dive deeper and adjust them.

MLOps is an analogy to (or, in some cases, an extension of) the well-known DevOps discipline, which transforms the way applications are developed (dev), deployed, and run (ops). In the ML world, it isn't code that's developed; it's a model built with certain data. When a model performs well during back testing, it's put in production to regularly forecast (ops). However, over time,

the parameters of the underlying data can change, or worse, the data quality can be influenced by technical errors in the process of data capture and preparation. Consequently, the accuracy of forecasts would drop. It's often critical to find out what happened and why it happened, as well as to react to such situations as fast as possible. This ability is what makes TIM unique. It brings the functionality to help in such situations, including obtaining warnings related to the data, quickly rebuilding models with the RTInstantML technology, and more.



TIP

TIM Studio can work with comma-separated value (CSV) files, as well as with data queried from SQL databases.

IN THIS CHAPTER

- » Deploying the Tangent Information Modeler in the cloud
- » Leveraging analytics and business intelligence tools
- » Using data integration tools
- » Bringing the Tangent Information Modeler to machine learning platforms
- » Extending the Tangent Information Modeler to the Internet of Things

Chapter 7

Using the Tangent Information Modeler on Other Platforms

As I explain in Chapter 6, all functionalities of the Tangent Information Modeler (TIM) Engine can be accessed through a single application programming interface (API). This API provides a single straightforward way of consuming the engine's capabilities, but also ensures an easy integration of these capabilities in various other platforms. This is an advantage for companies that want to leverage TIM's capabilities as built-in functionalities in their own platforms. It also provides opportunities for integration with countless other existing programs.

Building upon these opportunities, TIM can be accessed through a variety of tools and platforms. This is exciting because it means users don't have to switch tools every time they want a forecast or anomaly detection. In contrast, TIM's models, forecasts, and anomaly detections can be consumed directly in their platform of choice.

In this chapter, you learn about the growing ecosystem of TIM platforms offered by Tangent Works, allowing seamless integration of TIM's capabilities into users' familiar digital environments.

There are different types of platforms — including cloud, analytics and business intelligence (BI), data integration, machine learning (ML), and Internet of Things (IoT) platforms — each tailored to a specific segment of users. In striving to successfully deliver the most relevant functionalities to their user base, each of these platforms brings its own strengths. By integrating TIM with these different platforms, Tangent Works builds upon these strengths, focusing on the most appropriate of TIM's characteristics for each integration. Apart from looking at the set of integrations that is available to users, this chapter also gives an overview of the different types of platforms that TIM can be integrated with. This can help you decide which integrations best fit your needs.



TIP

The list of platforms TIM integrates with is always growing as Tangent Works improves existing integrations with additional functionalities and integrates with new platforms. For an up-to-date overview of TIM's platform integrations, take a look at www.tangent.works/tim-on-platforms and <https://docs.tangent.works/TIM-on-Platforms/Overview>.



REMEMBER

The ability to integrate TIM with different cloud, analytics and BI, data integration, ML, and IoT platforms provides many business benefits including the following:

- » AI, ML, and predictive analytics can be integrated in an existing environment.
- » Predictive analytics becomes easily accessible because it's available in platforms that are already being used in the organization instead of adding more tools to the mix.
- » It offers users broad availability, especially when it's available on multiple platforms.
- » It helps shorten time to market.
- » It allows users to easily tap into their data sources in the most suitable platform for their specific source.

Cloud Platforms

The TIM solution is a containerized application that can be deployed on any cloud infrastructure, including public clouds, such as Amazon Web Services (AWS), Microsoft Azure, Red Hat Cloud Suite, and private clouds.

Often, cloud platform vendors also have a rich set of data integration, orchestration, and presentation functionalities. A great example is Microsoft Azure. TIM is available in the Microsoft Azure Marketplace as a service and can easily be integrated with Azure Synapse, the Microsoft Power platform, and more.

Analytics and Business Intelligence Platforms

Analytics and BI platforms deliver analytic content development by enabling nontechnical users to execute analytic workflows from data access, ingestion, and preparation to interactive data visualization and analysis. They provide capabilities for analysis, as well as information delivery capabilities (such as reports and dashboards) and platform integration capabilities. TIM introduces predictive and prescriptive analytics capabilities on time-series data in an easy, accurate, fast, and explainable way, directly in these platforms:

» **Microsoft Excel:** TIM Forecasting is available as a Microsoft Excel add-in, providing an intuitive interface to the TIM Engine in Excel. The TIM Forecasting application for Excel supports real-time instant ML (RTInstantML), discussed in Chapter 5, allowing users to get direct forecasts based on the data in their Excel spreadsheets. Communicating with the TIM Engine from within Excel enables users to leverage the capabilities and familiarity of Excel, along with the powerful time-series insights that can be realized with TIM.

TIM in Excel increases the ease of use of TIM's RTInstantML capabilities for Excel users. For example, time stamps can be recognized in each of Excel's native date-time formats, the data-set range is automatically extracted, and additional comments or notes in the worksheet are automatically



TIP

ignored. Users can choose any variable as a target, and they can select which of the predictor variables should be included in the forecast. On top of that, the add-in provides users with the option for automatic extensive visualization, including the target variable with a forecast, prediction intervals, and a look into the importance of the included predictors.

The TIM Forecasting add-in can be added to Excel through Microsoft AppSource (<https://appsource.microsoft.com>).

- » **Microsoft Power Platform:** Microsoft Power Platform includes Power BI, Power Apps, and Power Automate. TIM integrates with Power Platform to support users end-to-end to gain insights and create business value from their time-series data.
- » **Qlik Sense:** Qlik Sense is a modern data analytics platform that helps users create flexible, interactive data visualizations. The TIM server-side extension (SSE) gives Qlik Sense users the benefit of TIM's forecasting capabilities without the need to leave the familiar environment of Qlik Sense. The TIM SSE integrates seamlessly with Qlik Sense's native way of working, providing various functions that can be used directly in Qlik Sense.

The TIM SSE for Qlik Sense supports RTInstantML, allowing users to get direct forecasts based on the data in their hypercube. This version also helps users gain deeper insights into their data, the models produced and used by RTInstantML, and the resulting forecasts by providing functionality to look into the features used in the calculation of the forecast. The SSE is set up to act as if it were a native Qlik Sense functionality, meaning that users can select a subset of the data set to train on, decide which variables to include, and indicate what the forecasting horizon should be, all within their dashboard.

Data Integration Platforms

A prerequisite for leveraging ML is data. TIM integrates with a range of data integration tools that specialize in gathering data. If you're working with large amounts of data, chances are, you're already using data integration platforms to help you handle the

integration and management of your data. Because the architecture of the TIM Engine makes it easy to integrate in any platform, you can enjoy the benefits that TIM offers from within your preferred data integration platform, including the following:

» **Azure Data Factory:** Azure Data Factory, part of the Microsoft Azure Synapse solution, is a service built for all data integration needs and skill levels, enabling users to integrate data silos. Users can easily construct extract, transform, load (ETL) and extract, load, transform (ELT) processes without code in an intuitive visual environment or write their own code and visually integrate a wide range of data sources. Azure Data Factory can also be used for the orchestration of forecasting processes through its seamless integration with TIM Studio.

When the preferred data-set and/or model-building definition is in place, the set of templates available from Tangent Works will guide you through the next steps. The combination of Azure Data Factory and TIM ensures that triggering data updates, forecasting, or model rebuilding is only a matter of a few clicks, regardless of whether you need these capabilities on a regular basis or in ad hoc scenarios.

» **Cloudera:** Cloudera delivers an enterprise data cloud for any data, anywhere, from the edge to artificial intelligence (AI). The combination of Cloudera and TIM provides an easy way to benefit from the scalable data management capabilities in Cloudera and deliver predictive analytics use cases with the augmented InstantML capabilities of TIM. The Cloudera and TIM InstantML integration offers a great time to market for your predictive and prescriptive analytics projects.

» **Snowflake:** Snowflake is a data platform built from the ground up for the cloud. It's designed with a patented architecture to be the centerpiece for data pipelines, data warehousing, data lakes, and data application development, as well as for building data exchanges to easily and securely share governed data. Snowflake offers great capabilities for data; TIM plugs into the Snowflake solution and delivers an integrated platform for augmented predictive analytics ML.

Many other data integration platforms exist, including Denodo, Trifacta, Databricks, and many others. The open architecture of TIM allows for easy integration into any platform.

Machine Learning Platforms

The time-series ML capabilities in the TIM solution are also available on various ML platforms. Users can make use of TIM's capabilities right next to any other ML endeavors they're working on, including the following:

» **Alteryx:** Alteryx Designer empowers data analysts by combining data preparation, data blending, and analytics; TIM expands this benefit with advanced time-series ML capabilities. The TIM Forecasting tool in Alteryx offers an easy-to-use, fast, and accurate solution for model generation and forecasting using explainable AI. The TIM Anomaly Detection Build Model tool and the TIM Anomaly Detection Detect tool together empower users to easily include TIM's model-building and anomaly-detection capabilities in their workflows. Business users, citizen data scientists, and data scientists benefit from the combination of the Alteryx functionalities and TIM's analytical capabilities, empowering them to create business value with data.



REMEMBER

ML for time-series data presents many obstacles, but it also contains enormous potential to gain deeper insights and deliver faster decisions. TIM focuses on equipping users to create business value through ML, instead of having their talent stuck on the technical aspects of complicated ML solutions.

The TIM Forecasting tool in Alteryx supports RTInstantML and allows customization of the forecasting models through adjustable advanced settings. The TIM Anomaly Detection Build Model tool also includes an interface to configure advanced settings. Both tools allow users to choose any variable as a target and select which of the additional predictor variables should be included in the forecast.

» **Azure Machine Learning:** TIM seamlessly integrates with the Azure Machine Learning solution. TIM offers augmented ML in Azure Data and ML pipelines. You can benefit from Azure MLOps (that is, DevOps for ML) by enabling data science and IT teams to collaborate and increase the pace of model development and deployment via monitoring, validation, and governance of ML models.

Internet of Things Platforms

TIM is architected to make efficient use of computational resources. Thanks to its design, TIM can be embedded into IoT devices or IoT platforms, equipping them with TIM's predictive and prescriptive analytics capabilities. This capability means users no longer need to centralize their data before starting to acquire insights. TIM can forecast expected values and detect anomalous values immediately after the data is measured. It's even possible to rebuild or retrain models on the edge (for example, in case of a structural change in the data, or when a certain sensor is faulty).

The IoT platforms TIM works with include the following:

- » **Azure IoT Central:** You can quickly build and deploy secure, scalable IoT applications using the comprehensive Azure IoT portfolio of managed and platform services like IoT Central and IoT Hub. You can benefit from TIM's augmented ML and deliver easy-to-use predictive analytics based on IoT data.
- » **Siemens:** Siemens MindSphere is an industrial IoT-as-a-service solution. Using advanced analytics and AI, MindSphere powers IoT solutions from the edge to the cloud with data from connected products, plants, and systems to optimize operations, create better-quality products, and deploy new business models. By adding TIM to the Siemens platform, augmented ML becomes available to many additional use cases.

IN THIS CHAPTER

- » Delivering key insights and business value with augmented analytics
- » Looking at the paradigm shift of the InstantML approach of the Tangent Information Modeler
- » Recognizing the speed, automation, accuracy and ease-of-use benefits of the Tangent Information Modeler

Chapter 8

Ten Ways to Get Value from the Tangent Information Modeler

In this chapter, I present ten ways you can get value for your organization from predictive analytics and the Tangent Information Modeler (TIM).

Driving Digital Transformation

Data holds the key to business insights that drive digital transformation initiatives in enterprises everywhere. Predictive analytics surfaces the information that business leaders and stakeholders need to make the right decisions at the right time.

Focusing on Business Value

A predictive analytics project isn't a math experiment. Avoid getting lost in the math and focus instead on the business value that the project will deliver. TIM's approach to time-series modeling

allows users to focus on the business insights that are contained in their data, while conveniently automating the handling of technical complexities under the hood.

Going Beyond Experimentation

Many machine learning (ML) projects never get out of the experimental stage. Data scientists get bogged down in manually building new models, and then training, testing, and tuning these models. This process can take days or weeks, during which time the business situation — and data structure — can change, often rendering the model obsolete. To avoid this trap, the projects need to be tightly coupled to the business problems they should solve. Many time-series problems are characterized by a need for fast and scalable insights. TIM's highly automated model-building capabilities are designed to attack these problems.

Getting Business Value through Augmented Analytics

Information is valuable but extracting it from raw data is often difficult and requires specialized expertise. This is particularly true in time-series analysis because of its highly dynamic nature. Many time-series use cases come down to forecasting and/or anomaly detection — areas in which TIM excels. TIM reduces the need for valuable resources (expertise, time, and money) and helps users leverage the insights hidden in their data to deliver real business value.

Approaching Time-Series ML in a New Way with InstantML

The TIM Engine creates models in a single step — from feature engineering to model building and deployment. This highly automated approach to time-series modeling is called InstantML. The high level of automation reduces the time needed for model building, as well as the engineering effort and mathematical expertise required.

Getting Results Fast

Traditional handcrafted ML models typically take days or weeks to build, train, test, and tune. AutoML reduces that time to hours or days, but it's compute-intensive and still requires extensive engineering support to build and deploy. InstantML delivers results within seconds or, at most, a few minutes with one-step model creation, while reducing the need for engineering support.

Automating the Model-Building Process

Many modeling tools require tedious, manual feature engineering that calls for a domain expert, such as a data scientist. This compute-intensive trial-and-error process is necessary to determine which combinations of input data will be relevant. This process can take days or weeks, especially as the number of combinations grows exponentially with the number of input variables and time intervals. On top of this, the algorithms used to model the resulting features often need parameter tuning before delivering optimized results. Tuning these parameters, either by hand or by use of AutoML, also requires expertise.

TIM automates the feature engineering process, analyzing the historical input data and determining which features are relevant given the use case, without the need for dedicated expertise. After the relevant features are determined, TIM builds an explainable model using these features and provides users with the desired forecast or anomaly detection.

Generating Accurate Models

Many ML models are built and tuned manually, requiring specialized skills, which are often limited to your data scientists. TIM's thorough automation allows the engine to create models with equivalent or better accuracy in a matter of minutes.

Explaining Model Insights

To solve time-series forecasting and anomaly detection problems effectively, business users need to understand the models and be able to explain them to stakeholders. This doesn't mean you need to be a data scientist, but you do need to have the confidence in your results to act on the information. Understanding the models and why they deliver a certain result greatly helps in creating this confidence.

TIM generates transparent, human-readable models that provide users with comprehensive insights into the models created and helps them measure the impact of predictors and features on target values — automatically and instantly.

Integrating in Your Existing Landscape Easily

TIM can be accessed across a broad ecosystem of tools and platforms including public clouds, analytics and business intelligence (BI) platforms, data integration platforms, ML platforms, and Internet of Things (IoT) devices. This makes it easier to leverage TIM's predictive analytics capabilities, because they're available in applications and platforms that are already familiar to your users.


TIM's architecture is optimized for seamless integration with existing databases, BI tools, and other enterprise applications. All of TIM's functionality is easily accessible through a representational state transfer (REST) application programming interface (API) that ensures the utmost deployment flexibility.

Notes

Notes

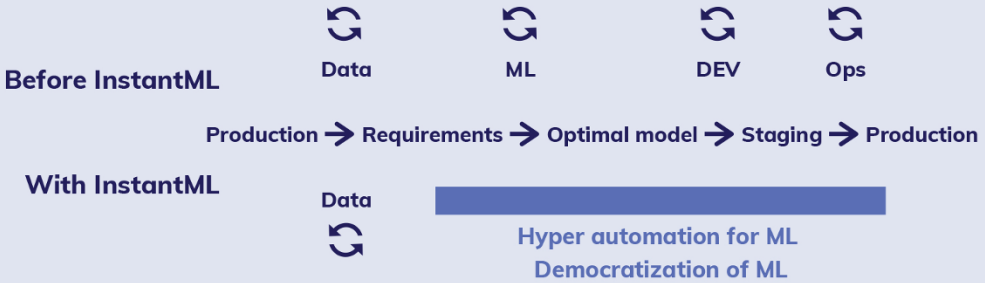
Notes

Notes



The Machine Learning **REVOLUTION**

InstantML, bringing predictive analytics to business users, getting business value out of time series data.



InstantML: fast, automated, easy to use, accurate and explainable Machine Learning



Tangent Works

Create predictive models in seconds for faster and better forecasting & anomaly detection

www.tangent.works - info@tangent.works

Corporate Headquarters
Oplobekstraat 6
1755 Gooik
Belgium

Tangent Works US
Tangent Works US
477 Route 10, Suite 208
Randolph, NJ 07869
USA

Tangent Works UK Ltd
The Innovation Centre
Sci-Tech Daresbury
Keckwick Lane, Daresbury
Cheshire - WA4 4FS

Tangent Works CEE
WESTEND TOWER
Dúbravská cesta 1793/2
841 04 Bratislava
Slovakia

Leverage time-series data in your business

Time-series data is everywhere. In industries from retail to finance and manufacturing to energy, companies try to use time-series data to deliver business value. But unfortunately, many machine learning projects never get past the experimentation stage as data scientists toil over handcrafted model building, testing, and tuning for days or weeks — far too long for businesses to leverage the information for real-time decisions. In this book, you learn about the opportunities and challenges of predictive analytics in time-series data, and how Tangent Works can help.

Inside...

- Deliver actionable insights in real time
- Explore business use cases for time-series data
- Recognize challenges in forecasting
- Integrate augmented machine learning in your platforms
- Accelerate time to value with InstantML
- Automate forecasting and anomaly detection on time-series data



Tangent Works

TIM

Lawrence Miller has worked in information technology in various industries for more than 25 years. He is the co-author of *CISSP For Dummies* and has written more than 180 *For Dummies* books on numerous technology and security topics.

Cover Image: Tangent Works

Go to **Dummies.com**[™]
for videos, step-by-step photos,
how-to articles, or to shop!

ISBN: 978-1-119-76753-4
Not For Resale

for
dummies[®]
A Wiley Brand



WILEY END USER LICENSE AGREEMENT

Go to www.wiley.com/go/eula to access Wiley's ebook EULA.