

Unlocking Solutions Superpowers with AI



Contents

Introduction

Vivun's AI Guiding Principles 3

Laying the Foundation for Artificial
Intelligence with a Solutions Ontology 5

Answering Fundamental Solutions
Questions 6

Delivering Amazing Buyer Experiences 13

Introduction

Why do Solutions teams need AI? B2B buying has changed for good. Now more than ever, buyers want in-depth, hands-on experience with your products before they sign a contract. They place their bets with solution experts who understand their business needs—and leading companies are growing their Solutions Organizations rapidly in response.

There is potential to capture and leverage the data and insights surfaced by the Solutions Org to improve visibility into individual deals, automate low value, repetitive tasks, and help companies keep product roadmaps aligned with market needs. Without a dedicated platform, this unique data set gets lost in the noise and goes unharnessed.

Vivun's Buyer Experience (BX) Platform is built to realize the full potential of this strategic function, and meet the expectations of today's sales-proof buyers. We do this with artificial intelligence and machine learning (AI/ML) to make Solutions data accessible across your entire company, and actionable for smarter, faster decisions and task automation.

What follows is an explanation of the various ways in which Vivun combines Solutions expertise with AI/ML to create uniquely intelligent recommendations on how Solutions leaders can unlock revenue, dominate their markets, and enrich their teams. We will cover the guiding principles that shape our approach to AI/ML, our Solutions Ontology, and how we use methods such as Bayesian inference, K-means clustering, natural language processing and generation in the Vivun platform.

Vivun's AI Guiding Principles

At Vivun, our approach to artificial intelligence is informed by several core tenets:

- **Immediate Value:** Solutions teams don't have the luxury of time when it comes to working deals, so we want our users to receive guidance from our AI from the very first day they start using the Vivun platform.
- **Continuous Learning:** As the market and competition continue to evolve, so should the platform's ability to make predictions and recommendations. Our machine learning models adapt to changing dynamics and new data submitted by Vivun customers over time.
- **Explainability:** Insights surfaced as a result of our AI/ML work need to be accessible and actionable to everyone in the company—with clear and concise explanations for numerical outputs (i.e. our Hero Score) and prescriptive advice on next steps to take, as opposed to a "black box" solution.

To meet all of these goals, Vivun's artificial intelligence combines an [expert system](#) and [machine learning](#) to predict what Solutions activities lead to successfully closed deals, and makes extensive use of [natural language processing](#) (NLP) to ensure those predictions are easily understood by our users.

Expert systems vs machine learning in artificial intelligence

Expert systems attempt to solve problems by applying prior knowledge that has been codified into a complex set of rules. These systems are particularly well suited to making predictions in situations where you have sparse data, but are [relatively inflexible](#)—expert systems do not learn from new observations or evolve over time unless they are modified by their creators, and their strict adherence to rules means that expert systems produce similar predictions for everyone that uses them, even with different inputs.

Many common machine learning algorithms, on the other hand, are characterized as "[stochastic](#)", indicating that their outcome has some notion of uncertainty or randomness. While stochastic learning approaches generally require a large amount of data to make reasonably accurate predictions, they can evolve over time as new observations are added to the dataset.

Machine learning can be subdivided into two main categories:

- **Supervised learning**, where input to the machine learning model is provided in the form of clearly labeled training data.
- **Unsupervised learning**, which asks the model to detect patterns without complete or labeled data.

Vivun builds software for people

One of our guiding principles for AI at Vivun is to make the insights surfaced by our platform actionable to every one of our users, as opposed to offering a "black box" solution. This is achieved in two main ways:

- **Explainability:** Any numerical output in the Vivun platform (e.g. the [Hero Score®](#)) is accompanied with a written explanation of what the score means, and what steps can be taken to further improve it, described via natural language generation (NLG).
- **Transparency:** In addition to providing an explanation of any numerical outputs, our platform shows users the structured inputs that go into the AI's analysis. Clearly indicating what inputs actually inform our AI-driven predictions helps Solutions leaders put our platform's recommendations into context when making decisions, or highlight issues with data quality that are preventing them from getting the most out of Vivun.

Laying the Foundation for Artificial Intelligence with a Solutions Ontology

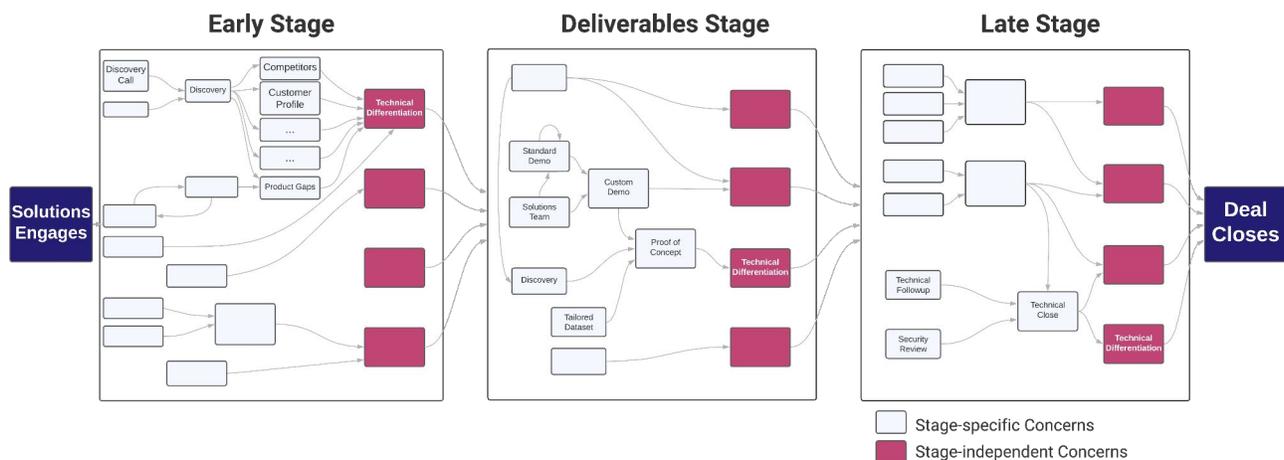
To build a model that could actually help Solutions orgs make reasonable predictions about the outcome of a deal from a technical perspective, we first needed to clearly define the world of Solutions Consulting and the different types of activity that lead to closed won opportunities by creating an ontology that captured the experience of Vivun's co-founders and the company's extensive research across the industry.

In the fields of information and computer science, **ontologies** show the properties of a subject area and how they are related to one another by defining a set of concepts and categories that represent that subject.

So how do we define success in bringing deals to closure? What concepts and categories might represent the world that Solutions Consultants inhabit? One might consider the various deliverables that Solutions teams are involved with at each step of an opportunity and how they are related to each other, such as technical qualification in initial calls, organizing a proof-of-concept as the deal proceeds, and providing additional assistance with technical evaluation in the final stages.

You could also look at a specific task—like discovery—and break down all of the different pieces of information that Solutions and their Sales partners typically gather about their customers. Recording these items and mapping them in relation to one another might eventually yield something resembling this:

A Simplified Representation of Vivun's Solutions Ontology



Vivun's **Solutions Ontology** is how we organize and describe the concepts and relationships that exist in the domain of Solutions Consulting. It is a structured representation of Solutions knowledge that draws upon the collective experience of Vivun's co-founders and the company's work with countless Solutions organizations, and is updated over time for our customers with data they provide from the specific deals they are working on.

Answering the Fundamental Questions of Solutions Teams

Hero by Vivun® leverages AI and machine learning to assist every Solutions professional in answering these fundamental questions about their work:

- Is this deal going to close? Why or why not? What will improve my odds?
- Am I spending my time on the right activities?
- What product gaps prevent us from closing deals, and how can we prove it?

**"How likely is this deal to close?
What will increase that likelihood?"**

Solutions consultants act as the technical conscience of every deal in the forecast. They keep their Sales counterparts honest by highlighting why a deal may be at risk from a technical perspective, and offer prescriptive advice on how to secure the technical win.

But even the best team members can be misled by gut feeling, or lack the ability to properly articulate why they feel a particular opportunity is at risk. At Vivun, the AI-driven Hero Score is our way of surfacing Solutions insights and improving the forecasting that teams already perform with sales tools. By facilitating the quantification and structuring of the technical conscience in every deal, our AI adds to the reasoning of the best members of your team—yielding more accurate forecasting, increased predictability, and alignment across the opportunity squad.

Improving the forecast with expert predictive capability

Predictions as to whether or not something will occur typically start with the foundation of what you already know, and are adjusted over time as you observe new things. The Hero Score follows this principle as well.

The Hero Score is driven initially by an expert system—an AI program that makes predictions by following prior knowledge codified into a set of rules. Because new users of Hero typically have a dataset on sales opportunities that is too sparse to make reliable predictions via machine learning, the Hero Score initially adheres to our ontology's expert assumptions about the domain of Solutions Consulting to make predictions about what activities and deliverables lead to a higher likelihood of Closed Won opportunities.

"Sales can sometimes fall victim to 'happy ears'—you want a repeatable, credible way to enter into that conversation with the PreSales perspective."

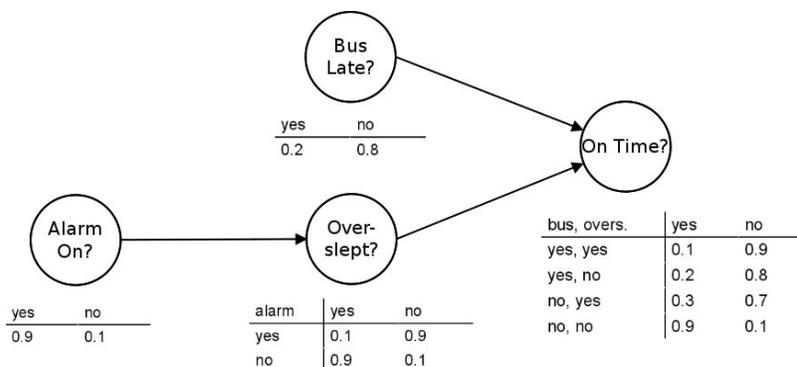
TOBY PENN

VP of PreSales, Recorded Future

Learning from your organization's unique attributes

Markets, competitors, and buyer behaviors inevitably change over time. Your forecasting and go-to-market strategy must adapt accordingly. As data is captured in Hero and deals are marked Closed Won or Closed Lost, the AI will begin to evolve with your business and apply new dynamics to the technical forecast. The learning system will autonomously track the nuances of your organization, and retain tribal knowledge about team members, blockers, and competitors.

To make more dynamic predictions on how likely a given deal is to close, our platform starts relying less on the expert system and bases more of its recommendations on the output of a probabilistic graphical model, specifically a [Bayesian network](#). Bayesian networks are ideal for taking events that occurred and modeling probabilistic relationships between them, for example the relationship between sleep and catching a bus—you could create a Bayesian network that computes the probability that you might miss the morning bus, given that you overslept your alarm.



Source: '[A Tutorial on Learning with Bayesian Networks](#)'.

A Bayesian network applied to Solutions orgs might ask the following: given a certain set of concerns (demo, proof-of-concept, product gaps, relationships), how likely is it that this opportunity will result in a Closed Won deal?

Using the concerns captured on every opportunity tracked in Hero as inputs to the model (present competitors, technical differentiators, opportunity gaps, deal momentum, and the opportunity team's track record), our Bayesian network computes the probability that a deal will close, and outputs the result in the form of the Hero Score.

An AI-driven prescription for healthy deals

Solutions insights surfaced in the Hero Score must be as accessible as possible, so we also use [natural language generation](#) (NLG) to provide a human-readable explanation of what the Hero Score means, as well as prescriptions on how to mitigate opportunity risk, improve data hygiene, and improve the likelihood of winning the opportunity based on deals that had similar attributes.

While charts and numbers can look visually stunning, a written explanation makes it straightforward for anyone using Hero to immediately understand what the Hero Score means for an opportunity they are working, and what steps are needed to win the deal—not just individuals who have a formal education in statistics or data science.

Working the right deals, the right way, every time

By synthesizing Solutions insights on every opportunity into a comprehensive explanation of how to secure the technical win on a deal, Vivun enables your team to work the right deals, improve forecasting accuracy, and increase deal velocity. Since the model behind the Hero Score learns from its experience about a particular organization, it retains the collective knowledge of the team, behaving as a team member that remembers everything—while never leaving your company.



"With Vivun, my PreSales leaders show me exactly where they are impacting revenue. This enables us to adapt our strategies to improve win rates, move faster, and close more business."

ED CALNAN
CRO

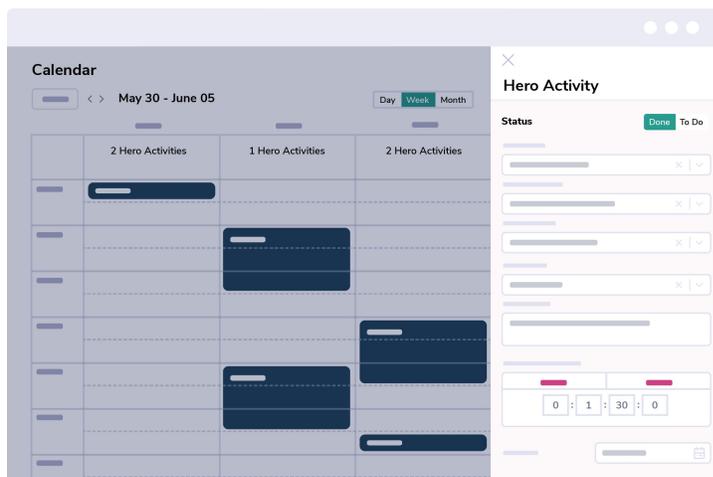


"Where is Solutions spending its time? Are we focused on the right things?"

Activity tracking is time-consuming and hard to prioritize because it does not directly drive revenue, but Solutions team members care deeply about the ability to capture and analyze their work. For individual contributors (ICs), efficiently tracking deliverables and activities is a great way to get credit and acknowledgement for their work on deals. For leaders, activity capture yields crucial data on how the team is spending its time—revealing where adjustments can be made to drive repeatable and efficient technical wins, or whether additional budget and headcount for the team is justified.

Integrating activity tracking where Solutions lives—their calendars

A leader's ability to improve their team's performance hinges on their ability to make capturing team member activities as painless as possible. One way to automate activity tracking is by mapping each Solutions team member's calendar events to the opportunities they're working on. Like many other customer-facing departments, Solutions teams' schedules are filled with external meetings—and the text and time blocks of calendar invites can help answer questions about where a team is spending its time, and how much effort was needed.



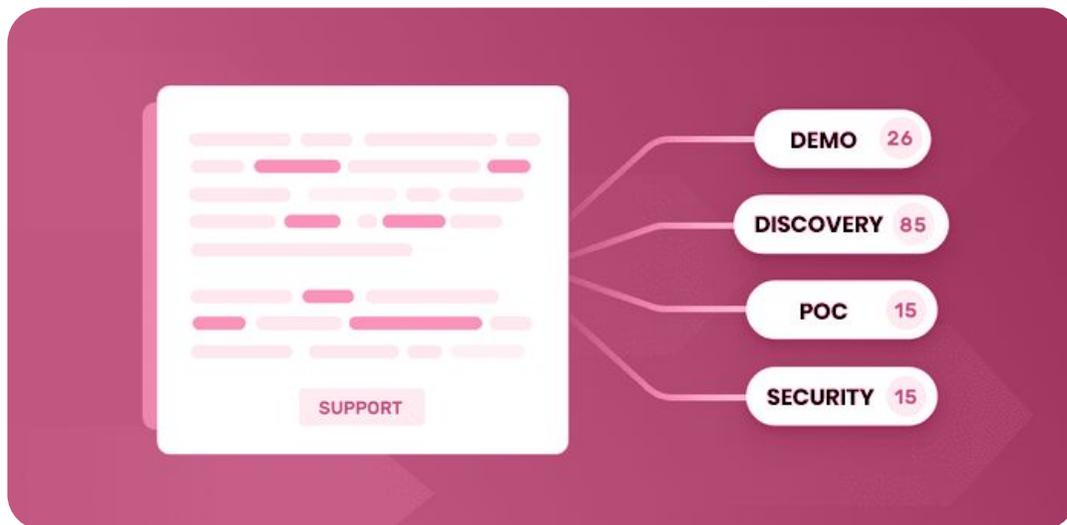
While it's fairly simple to look at a meeting invite and determine that it was an hour long, examining the language in that meeting invite to classify how that time was spent is less straightforward.

Furthermore, calendar events and activities exist mostly as words, rather than as numbers, so additional processing is needed before any kind of quantitative analysis can be performed. This step of converting text into numerical representation is referred to as [feature extraction or text vectorization](#) in the field of natural language processing.

Bringing automation and intelligence to activity tracking

In order to predict which calendar invites map to which deals, calendar integrations in Hero utilize [TF-IDF \(term frequency—inverse document frequency\)](#), which is a method that generates scores for how relevant a given word is in a set of documents by multiplying:

- **Term Frequency**, how many times a words appears in a given document
- **Inverse Document Frequency**, how many times that word appears across a set of documents



TF-IDF is especially useful for matching calendar activity to opportunities and deliverables because most calendar invites contain a set of words that is common to all of them, and then words and phrases that are much more rare.

Scores generated via TF-IDF are used to provide recommendations to the user on how each of their calendar events should be associated with specific opportunities they're tracking in Hero. Predicting the relevance of meetings to opportunities reduces the time that Solutions team members need to spend on administrative catching up, while helping Solutions leaders quickly identify whether valuable time is being spent on the right things.

"What product gaps are preventing us from closing deals? How can we make our case to Product?"

As technical liaisons embedded in the field, the Solutions Org is constantly evaluating product-buyer fit and uncovering new customer use cases. These insights are immediate, actionable, and can directly shape the product roadmap to drive growth. However, it's often the case that feature requests are driven by the largest deal in the forecast, the loudest set of voices in the room, or worst—completely unacknowledged due to a lack of data.

Even in organizations where collaboration between Product and the field occurs to capture customer feedback, the process is often manually intensive and incredibly noisy. Sorting through a mountain of tickets filled with duplicate requests from disparate systems to find what is truly important results in missed opportunities, and a lack of alignment about what matters most on the roadmap.

Seamlessly capturing and aggregating real-time product feedback

At Vivun, we automate the intake of feature request data, as well as the aggregation and analysis of that data into discrete Product Gaps that can be assessed and acted on by Product teams.

Whenever a new gap is recorded on an opportunity in Hero, it will be automatically matched with existing product gaps via natural language processing, preventing Product teams from being overwhelmed with duplicate feature requests.

As the total amount of data on Product Gaps grows, the system will automatically identify compelling patterns about your product and what buyers are asking for, such as the frequency of a given feature request, common attributes among customers who ask for a particular feature, or the overall impact to revenue.

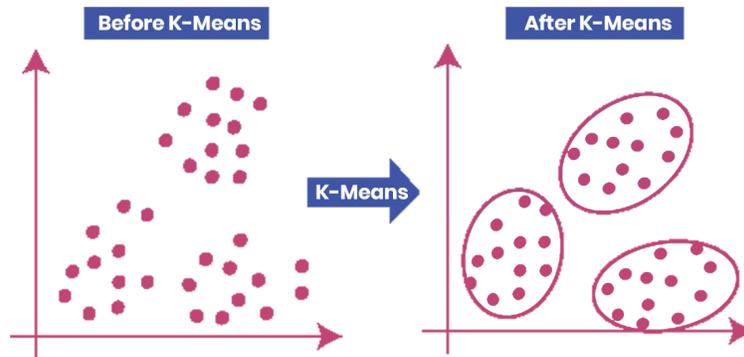
Match %	Template Name	Product Line	Percent Complete	Action
36%	Note Name Inherited from Template	Hero	0%	SELECT
24%	Notes	Hero	100%	SELECT
24%	Default Templates & Attachments			SELECT

"With the product intelligence we've derived from Vivun, we've been able to close both mid-market as well as enterprise deals faster, because we weren't chasing the wrong features. I can tie several hundred thousand dollars in ROI to our use of Vivun."

NICK DURKIN

Field CTO, Harness

After mapping data inputs to known product gaps, we use a [K-means clustering algorithm](#) to detect underlying patterns in the product feedback being submitted via Hero, so that Product teams get rich context on new use cases, growth opportunities, and customer goals.



The ability to sort a large amount of product feedback into several distinct buckets does not just eliminate duplicate feedback. It also assists in uncovering trends that you may not have been previously aware of. As an unsupervised machine learning algorithm, K-means clustering is ideal for processing data that is not clearly labeled—Solutions teams often have a universal set of terms to describe their deliverables and activities on a given opportunity (i.e. demo, proof of concept, technical evaluation), but product feedback and information on customer use cases generally uses much more varied language.

Over the long term, Hero users can leverage this data to evolve the Solutions Ontology and how heavily specific feature requests impact the Hero Score, and gain the ability to answer questions like how long a feature request went unresolved and what the impact to revenue was during that period, or what competitors were most relevant at the time a particular product gap was first logged, and how that compares with more recent data.

Determining what to build, and measuring impact with data

To help with prioritization and justify product investments, Vivun provides business metrics for Product Gaps, tracking the frequency, status, and value of the feature requests you receive from prospects and customers. Because product gap data is tracked and visualized over time, the entire team can see what is needed to close open opportunities and drive more revenue.

Furthermore, Hero reduces the amount of work needed to keep Product connected to the field in multiple ways. Teams in the field can easily reference relevant documentation for Product Gaps. Tight integrations with the solutions that Product teams live in—GitHub, Jira, and Azure DevOps—mean that the field is automatically notified when the company closes product gaps that are negatively affecting revenue, and can drive users to new features to revive stalled or lost deals.

Deliver Incredible Buyer Experiences

Solutions teams are essential to delivering the experience today's technology buyers demand. Providing Solutions leaders with the AI-powered ability to automate routine tasks and turn data into actionable insights lets them manage teams more effectively, align sales with product, and find ways to turn deal wins into repeatable best practices.

Vivun is the leader in Buyer Experience (BX) software. Its AI-powered platform enables B2B businesses to reach today's sales-proof buyer, differentiate from the competition, and accelerate revenue. With the Solutions Organization at the forefront, Sales and Product align, workflows scale, and more deals are won with transparency and trust.

The world's leading companies power incredible buyer experiences with Vivun.
Learn more at www.vivun.com.

zoom

 snowflake

okta

 SecureWorks

coinbase

vivun.