Capturing your invisible market

Volume I, GTM Intelligence

GoodFit

Every revenue leader I talk to is being asked to deliver more growth from channels that are getting harder.

Everyone is targeting the same set of accounts with the same set of channels. The result is spiraling CAC and prospects that never hear what you have to say.

The industry's answer has been to throw tools at the problem. There are more sales platforms, more AI SDRs, more intent signals and orchestration layers than at any point in the history of B2B. And yet none of it is working. Reply rates are at historic lows. Pipeline is harder to build than it has ever been.

None of it is working because the tools are not the problem.

We believe that the fundamental problem in sales today is that revenue leaders don't have clarity on what accounts to target or how to allocate resources against them.

This is the first volume of GTM Intelligence, a series where we'll share what we're learning from working with revenue teams to define, source & grade their market.

Volume I is our first and most foundational piece. On average, B2B companies are missing 60% of the market they could sell to, while 70% of the accounts in their CRM don't match who actually buys their product at all. We'll show you why that gap exists, what it's costing you, and what to do about it.

If you've ever had a nagging feeling that your team might not actually be targeting the right accounts... or that you're missing visibility into large swathes of the market... this one is for you.

— Harrison Rose, CEO, GoodFit

Charlie Lynch
Sales Director at Merge

I cannot imagine joining a new company as a sales leader and not starting by mapping every company in my market.

Capturing every single account you can sell to is the critical first step in building a revenue engine.

Every decision you make — what accounts to prioritise, what go-to-market strategies to employ, how many reps to hire — depends on having a clear, qualified view of every account you can sell to.

But our research shows that, on average, B2B companies are missing 60% of their market. There is a fundamental mismatch between who companies think they sell to, who they have sold to in the past, the accounts they get from their data provider, and who actually buys their product.

ICP definitions have been built off proxy filters that fail to capture the nuance of who actually buys products; the result is an invisible, uncaptured market of buyers that live outside standard filters.

Your invisible market caps your potential revenue before you even get started. Your TAM model, your territory plans, your hiring forecast... are all built on a foundation that's missing more than half the picture. There are three reasons why your invisible market exists. Here they are, and how to combat them.

1
Your ICP is probably wrong
Page 6

2
The accounts you're looking for can't be found... with filters
Page 15

3
Your market changes every day
Page 24

Your ICP is probably wrong

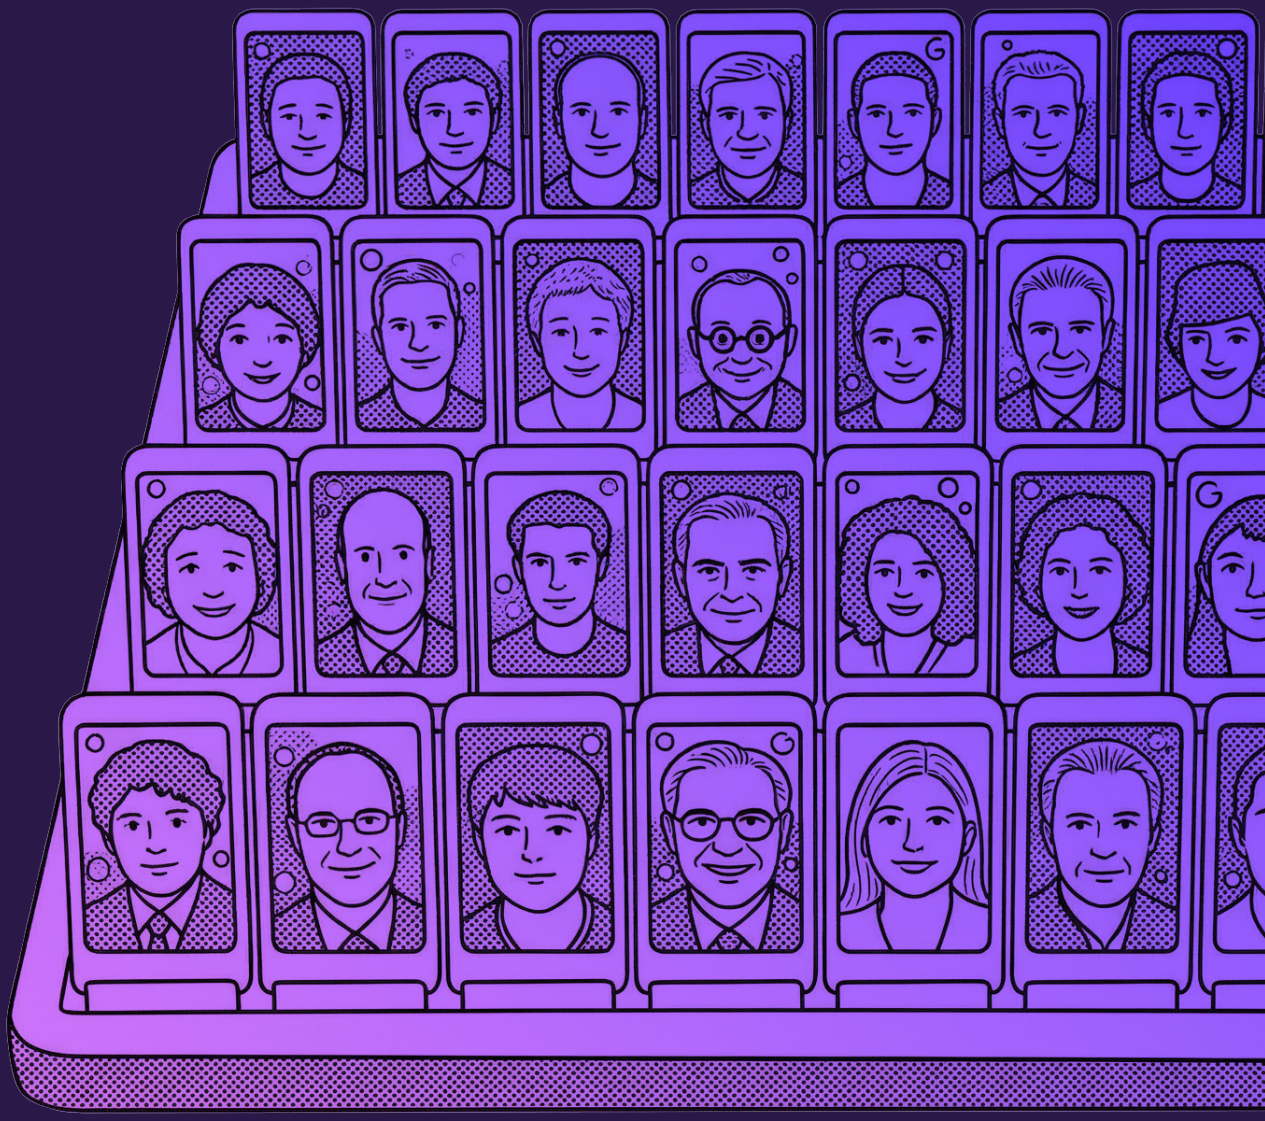

For years, mapping your market meant going to a data provider, accessing the same dataset that every other company could see, and filtering by a combination of employee count, industry and geography.

These three filters were the standard because they were the only data points available. They were simply the only data points that providers could reliably track with enough coverage to know them for every company in their database.

The problem is that these filters were never designed to identify your actual buyers. They were the best approximation available given the data that existed. And now, companies have been using these proxies for so long that many have forgotten what their actual ICP was in the first place.

Ask a CEO who they sell to, and they'll often say something like: "50 to 500 employee technology companies in North America and Europe." But that is not actually an ICP — it's a proxy for one. Their actual ICP might be: "B2B SaaS companies with an established inbound sales motion." The legacy filters couldn't capture that, so the proxy became the truth.

Percentage of accounts in the average B2B company's CRM that do not match their actual ICP

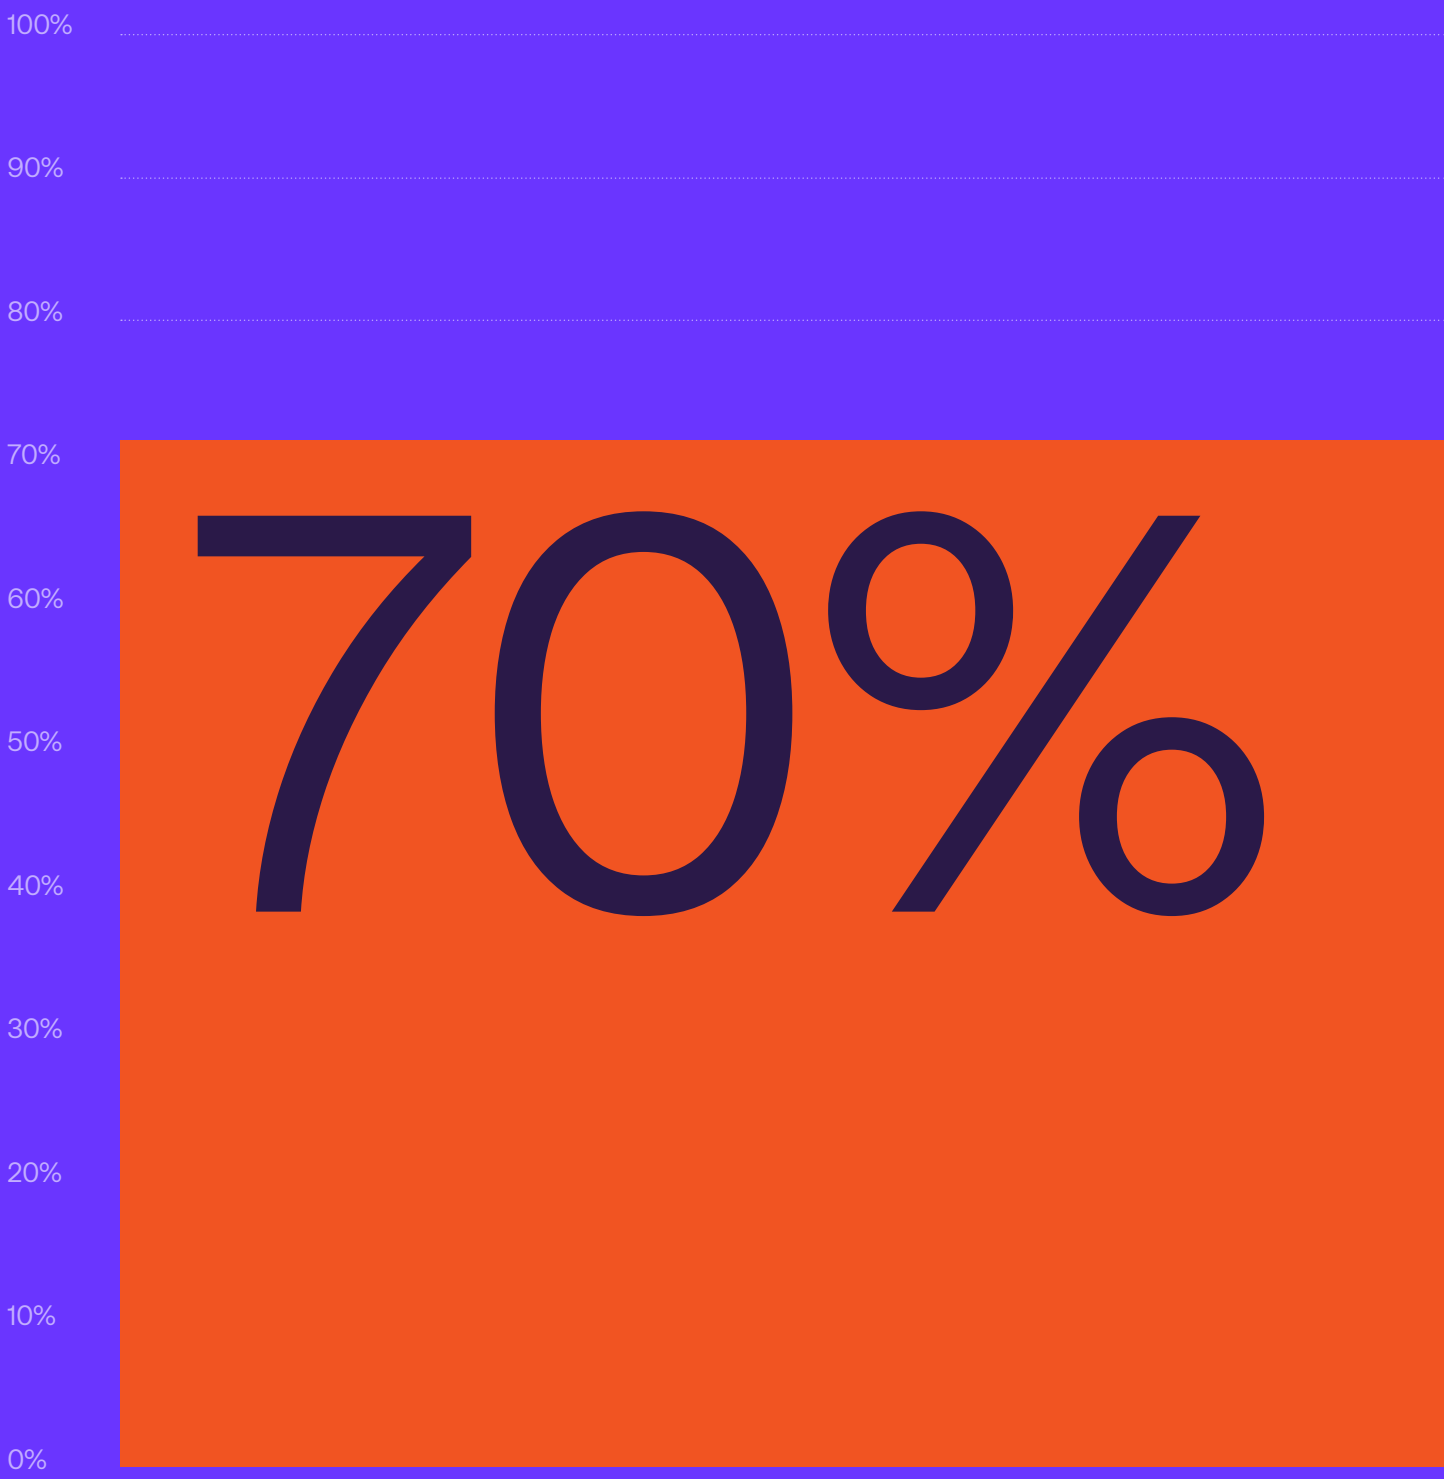

70% of accounts in the average B2B company's CRM that do not match their actual ICP. These are the non-ICP accounts that your proxy filters capture by casting such a broad net, or that get added to your CRM by self prospecting and random inbound leads, or that once fit your ICP filters but have slipped out. The result is that 70% of the average B2B CRM is junk — accounts that will never become a deal, but will create noise and consume rep time, marketing budget and management attention. Literally every dollar and every minute spent on these accounts is wasted; they will never buy your product.

Capturing your real buyer

Most ICPs are built across three dimensions. These dimensions stay the same; what changes is the data within.

Company size. This has traditionally been proxied through employee count. But employee count is a blunt instrument — it is often self-reported and inaccurate, and it doesn't tell you what you actually need to know. Depending on your product, "size" might be better captured by the size of specific teams (how big is the sales team? the finance team?), revenue, number of customers, or even web traffic volume. A 200-person company with a 3-person sales team and a 200-person company with a 40-person sales team are completely different prospects if you're selling a per-seat sales tool.

Geography. This has traditionally been proxied through a single "country" field. But geography is more nuanced than HQ location. It includes where a company's employees actually sit, where their offices are, where their customers are, and where they're hiring. A company headquartered in London with 80% of employees outside the UK is a very different prospect to one headquartered in London with 80% of employees inside the UK.

Offering-to-business fit. This is where legacy data has failed most completely. Historically, "does this type of business fit our product?" was proxied through industry. But industry is a spectacularly poor proxy for whether a company needs your product. What you actually need to know is something about how the business operates: do they have an inbound sales motion? Do they process personal data? Do they hire across multiple countries? Do they have a particular tech stack?

Some companies are lucky enough to have clean, binary criteria: "we sell to every eCommerce retailer that uses Shopify." If that's you, your ICP translates directly into filters. But most companies have a qualitative ICP: "Companies with an established inbound sales motion," or "Data-centric businesses that process personal data." If that's you, the challenge is finding the right proxies to capture that ICP at scale.

A 200 person company with a 3 person sales team and a 200 person company with a 40 person sales team are completely different prospects if you're selling a per-seat sales tool.

A company headquartered in London with 80% of employees based outside the UK is a very different prospect to one headquartered in London with 80% of employees inside the UK.

Chili Piper

Chili Piper is a demand conversion platform that turns inbound website interest into qualified pipeline, through AI agents that engage, qualify, route, schedule and re-engage.

Chili Piper sells to B2B companies in North America and EMEA with a large inbound sales motion. Here's how they captured that market with legacy vs modern filters:

Legacy filters

50–500 employees, industry exclusions, North America + EMEA.

300K

accounts

Modern filters

Region is North America and EMEA

and **Sales employee count** is greater than 10

and **Total website visits** is greater than 50,000

and **Inbound CTA on website** is True

13,000

accounts

Playroll

Playroll is a global employment platform that helps companies hire, pay, and manage workers anywhere across EOR, contractors, and in-country payroll through a single unified system.

Playroll sells to any company hiring employees outside their HQ location. Here's how they captured that market with legacy vs modern filters:

Legacy filters

Global companies with 10-5,000 employees.

2.5M

accounts

Modern filters

Region is **Global**

and **Hiring** **+2** **roles outside their office location**

and **Team in** **+7** **countries**

and **People in Global Payroll** is greater than **1**

160,000

accounts

You don’t have to figure out your ICP manually

One of the most common mistakes in market mapping is trying to define your ICP qualitatively. You can find thought leaders on LinkedIn who proudly describe how they sat down with a company's best prospectors and sellers who "know in their gut" what makes a qualified account, then matched that with sales call transcripts, won-loss reasons and product team workshops. It sounds rigorous... but there's a better way.

Pull a list of qualified accounts (such as your customers), and a list of unqualified accounts (ones you know are not buyers). Enrich them with a deep set of company-level attributes, and run an analysis on what attributes they share. The data will tell you what your ICP actually is — often revealing criteria you hadn't considered and disproving gut feel assumptions.

If you already have a set of proxies you use to define your market, you can test them the same way: pull your customer list, check how many of your existing customers actually match your current criteria, and see what falls through the gaps.

Problem 2

The accounts you're looking for can't be found... with filters

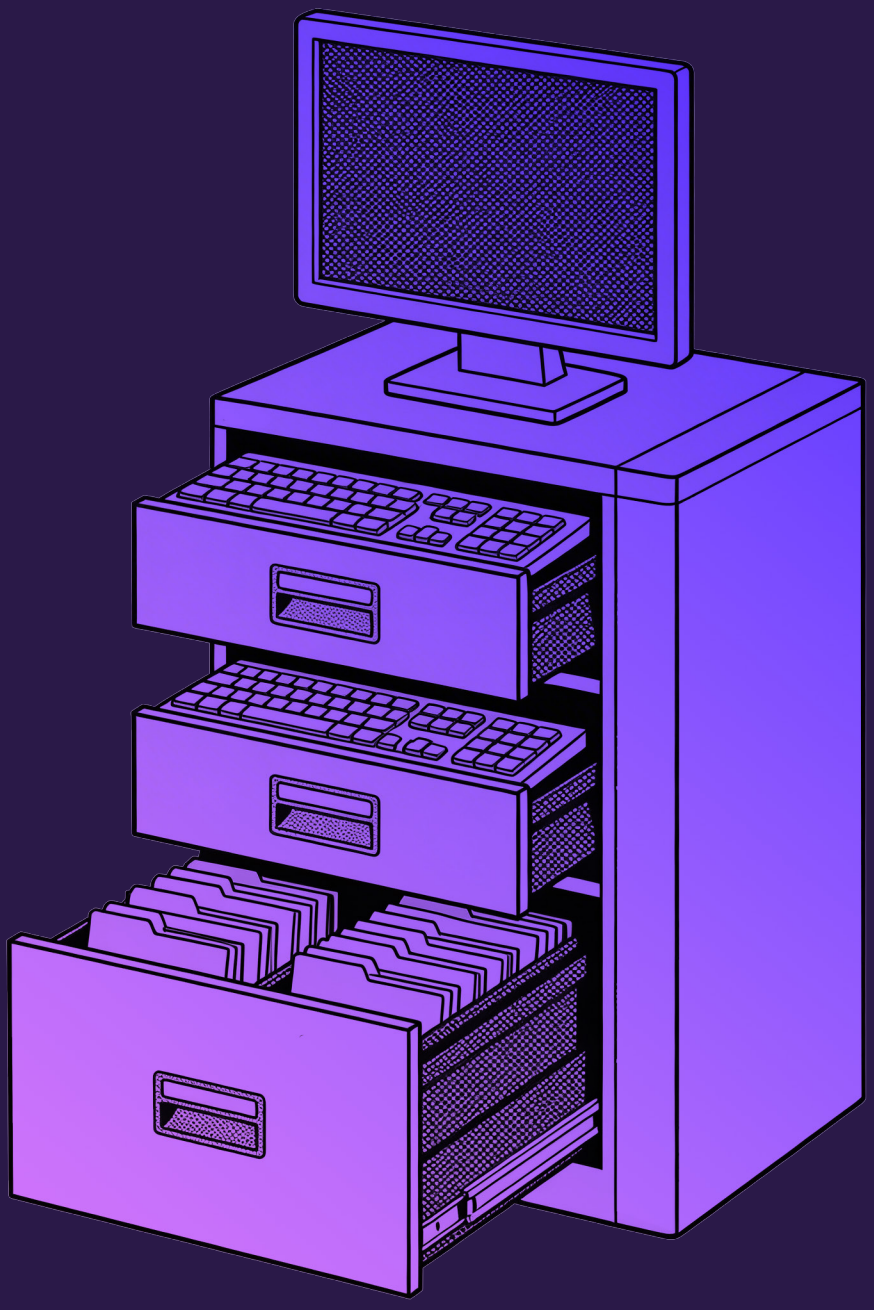

Average B2B company's actual ICP missing from their CRM

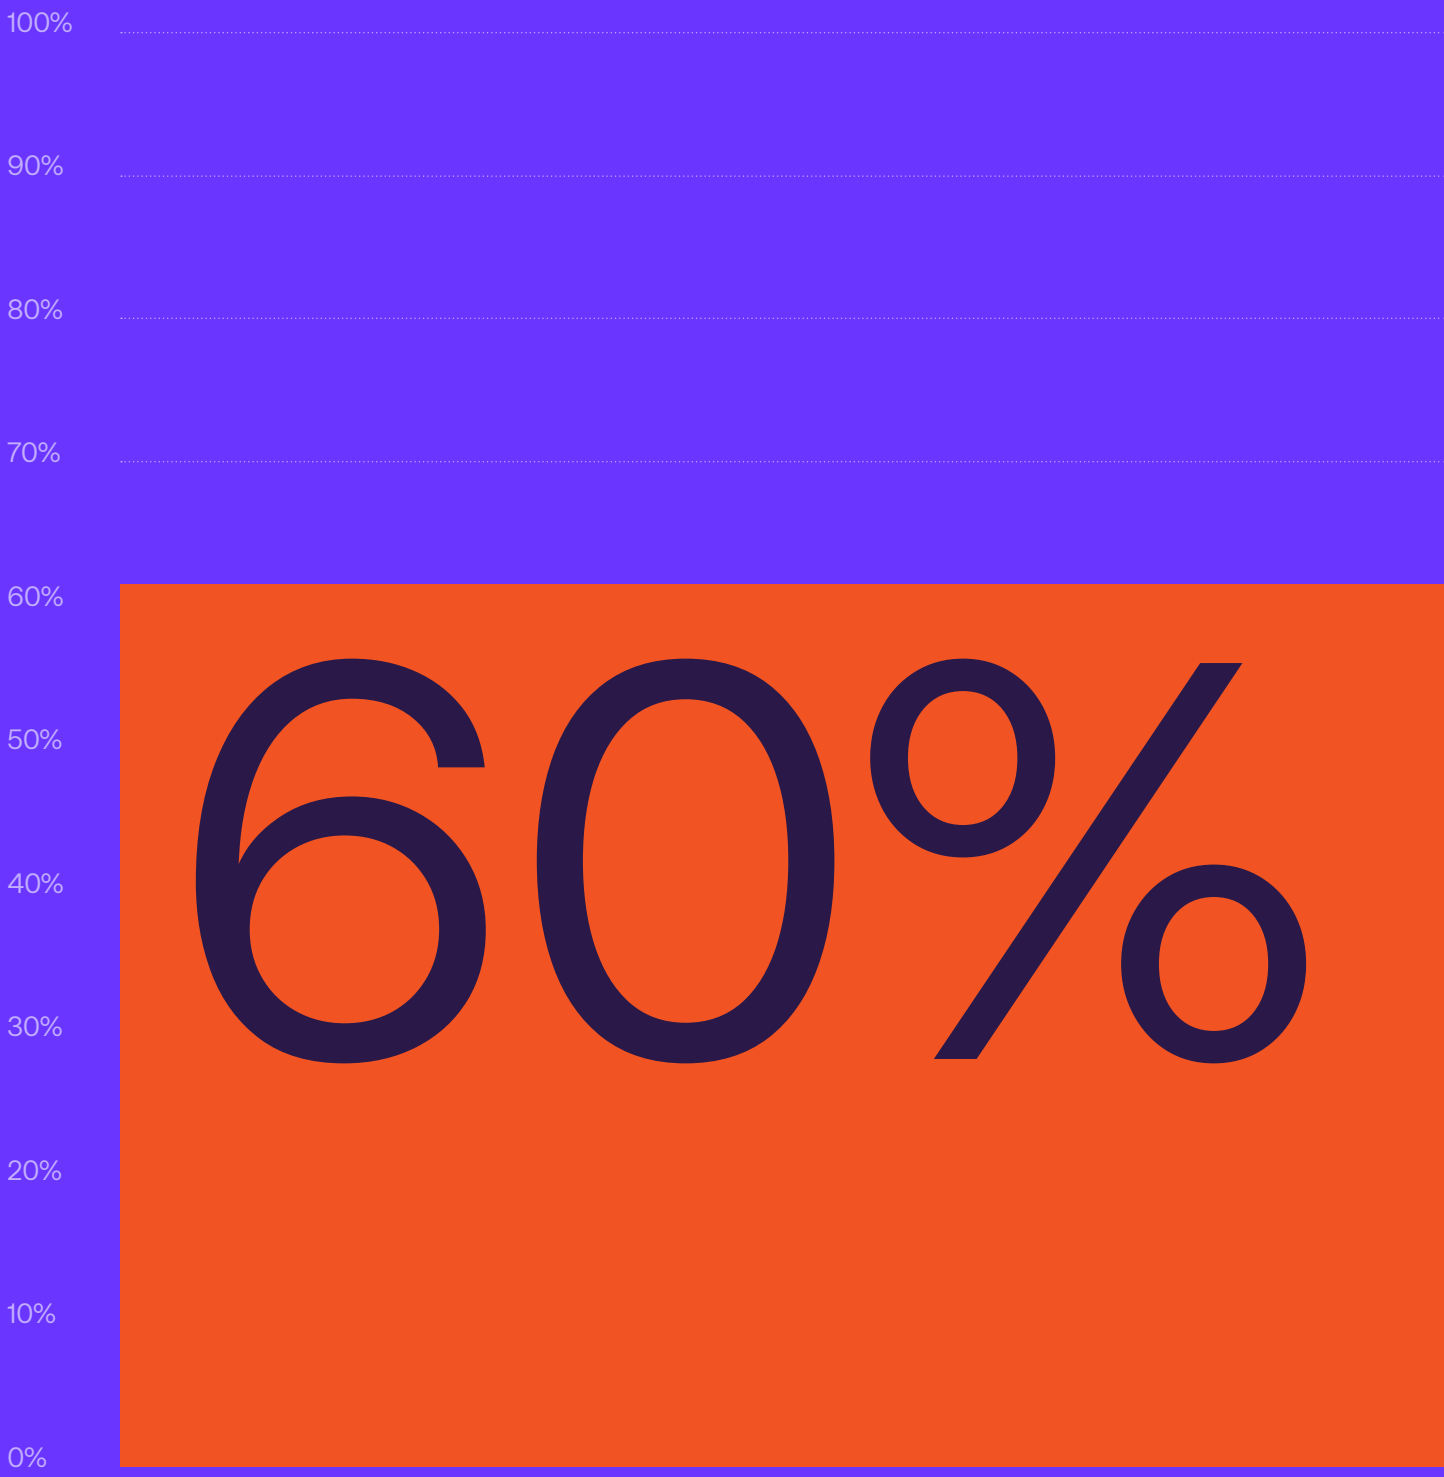

Once you've properly defined your market & worked out the filters that define your ICP, you can run into a second problem: sometimes your ICP can't be defined by filters at all.

Better data solves most of the "proxy problem." If your ICP is "companies with a large inbound sales motion," smarter filters get you a long way: you can look for 10+ person sales teams, 50k+ website visits, a CRM like HubSpot or Salesforce. These are far better proxies than "50–500 employees in Technology," and for many companies, they're precise enough to map a market with confidence.

But sometimes there's a far more elegant solution to the proxy problem. "Does this company have an inbound sales motion?" is a question a human could answer in thirty seconds by visiting the company's website and looking for an inbound call-to-action, like "Book a demo" or "Talk to Sales." You could try to approximate that through filters — reading job titles for signs of an inbound team, checking for marketing automation tools in their tech stack — but these are still proxies for a question that a rep could answer at a glance. The problem is that this data doesn't exist in any structured database. And you can't send a human to check 500,000 websites.

Classification models outperform LLMs when making data judgements

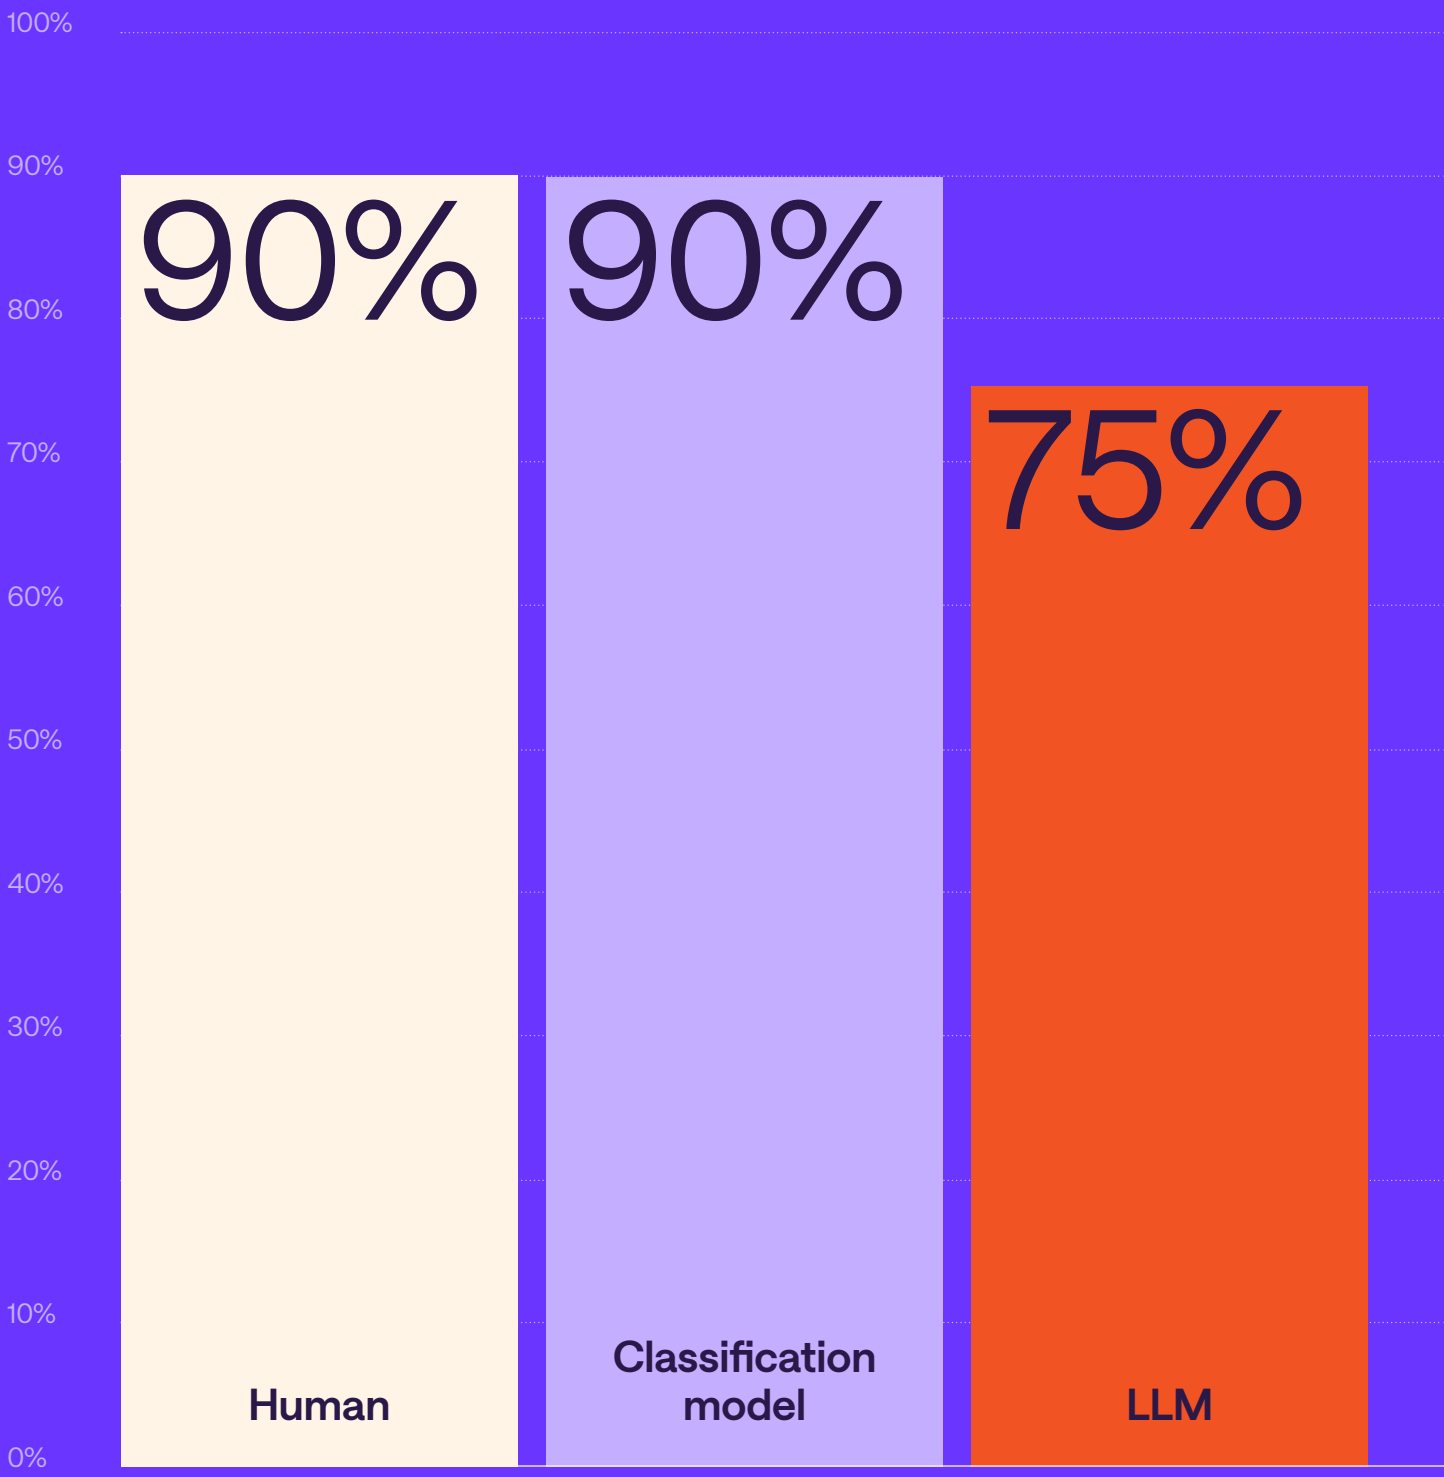

Using classification models to capture hidden accounts

When trying to answer questions that a human could answer, the natural instinct is to turn to an LLM. You could connect an LLM to a web scraping agent, give it a company's website and ask: "does this company have an inbound sales motion?" It sounds like the obvious solution, and for a handful of accounts, it might work well enough. But when you need to run that question reliably across your entire market, the cracks begin to show. LLMs are linguistic tools; they read text and interpret it. With grey-area examples, they struggle. At best, LLMs top out at around 85% accuracy on most classification tasks. More commonly they sit around 75%. Think about that in practice: you assign 100 leads based on an LLM classification, and 15 to 25 of them are wrong. If you can't get your lead assignment to at least 85% accuracy, your reps are going to lose confidence in the leads you give them.

Classification models work differently. Rather than reasoning through language, they learn the underlying quantitative patterns across thousands of labelled company examples. They are trained on a single, well-defined question — "does this company have an inbound sales motion?" — and they learn what the statistical signals of a yes or a no actually look like in the data. The result is that classification models consistently achieve human-level accuracy: 90% or higher.

The training process is a simple loop. You start with a set of labelled examples — say, 20 companies that clearly have an inbound motion and 20 that clearly don't. The model classifies a new batch, human experts review and correct the outputs, and the model retrains on progressively harder examples. This repeats until the model reaches or exceeds the accuracy of a human reviewer.

Classification models are not always the right tool. They are higher cost than filters, so the decision to use one should be deliberate. The logic is straightforward: if the criteria has a high impact on who is in or out of your market, and the information needed to assess it is publicly available at scale, a classification model is worth building. If it only affects a small percentage of accounts, or the information isn't reliably available online, a filter or manual check is more practical.

But when the fit is right, classification models let you move beyond proxies entirely. Instead of approximating your ICP through filters, you can answer the actual question — at scale, with human-level accuracy, across every account in your market.

The company and contact data market is optimised for volume rather than quality. That creates a problem for businesses like ours with a tightly defined ICP, especially when key ICP traits are difficult to identify at scale.

Mitchell Omer, Director of Revenue

TrustKeith is a privacy management & GDPR compliance solution. They sell to global SMBs that process lots of personal or sensitive data.

"Does this company process personal or sensitive data?" is not a field in any database.

Before deploying a classification model, their reps were spending over two days a week manually qualifying lists - visiting websites one by one, trying to answer that question themselves.

After deploying a model trained on exactly this question, they increased their total qualified accounts by 170%, with no human touch required.

TrustKeith's reps got two days a week back and their market got bigger.

Sourcing new accounts vs. enriching the ones you already have

There is a fundamental distinction that most revenue teams overlook: the difference between enriching accounts you already know about, and sourcing accounts you've never seen.

Most tools in the market — traditional data providers, enrichment platforms, scraping workflows - start from a list of companies you already have and add data to them. That's enrichment. It makes your existing accounts better, but it can't show you the accounts that aren't in your CRM at all.

Sourcing is different. It means building a comprehensive dataset of every company that matches your buyer criteria, regardless of whether you've heard of them. It requires a different kind of infrastructure: the ability to start from a universe of companies, apply your specific ICP criteria (including exotic signals that don't exist in any structured database), and return every account that qualifies.

This is why the 60% gap exists. It's not that companies aren't investing in data. It's that they're enriching what they can see, when the real opportunity is in sourcing what they can't.

Our TAM was basically every account in Apollo.

Charlie Lynch, Sales Director

Merge is the leading provider of agentic tools and customer-facing integrations, serving companies including U.S. Bank, OpenAI, BambooHR and more.

Merge Unified's ICP isn't defined by industry, employee count, or geography. It's defined by a signal that doesn't exist in any database: "Does this company offer integrations to its users?"

But "software company that offers integrations" isn't a filter in traditional data providers like ZoomInfo or Apollo. What those tools do offer is a definition of "software company" that is vastly broader than what Merge actually needed.

"Integrations offered" doesn't exist as a data point in any structured database; the only way to find it is to look at a company's actual website. But you can't scrape every website in the world. You need to start with a list of the right companies.

And Merge's problem was that they didn't have that list either: no data provider could reliably identify which companies in their database were software companies in the way Merge defined "software company."

The Merge team solved this problem in two steps. First, by finding a proprietary database that had accurately tagged 100,000 software companies (plug: it was GoodFit). Then, by scraping those companies' websites to detect which of them listed integrations that Merge offers (second plug: we handled that for them, too).

In doing so, the Merge team bypassed the proxy problem entirely: instead of using traditional data to ask, "is this the sort of company that would need to have integrations?", they simply asked the real question: "does this company offer integrations that match the Merge offering?".

Your market changes every day

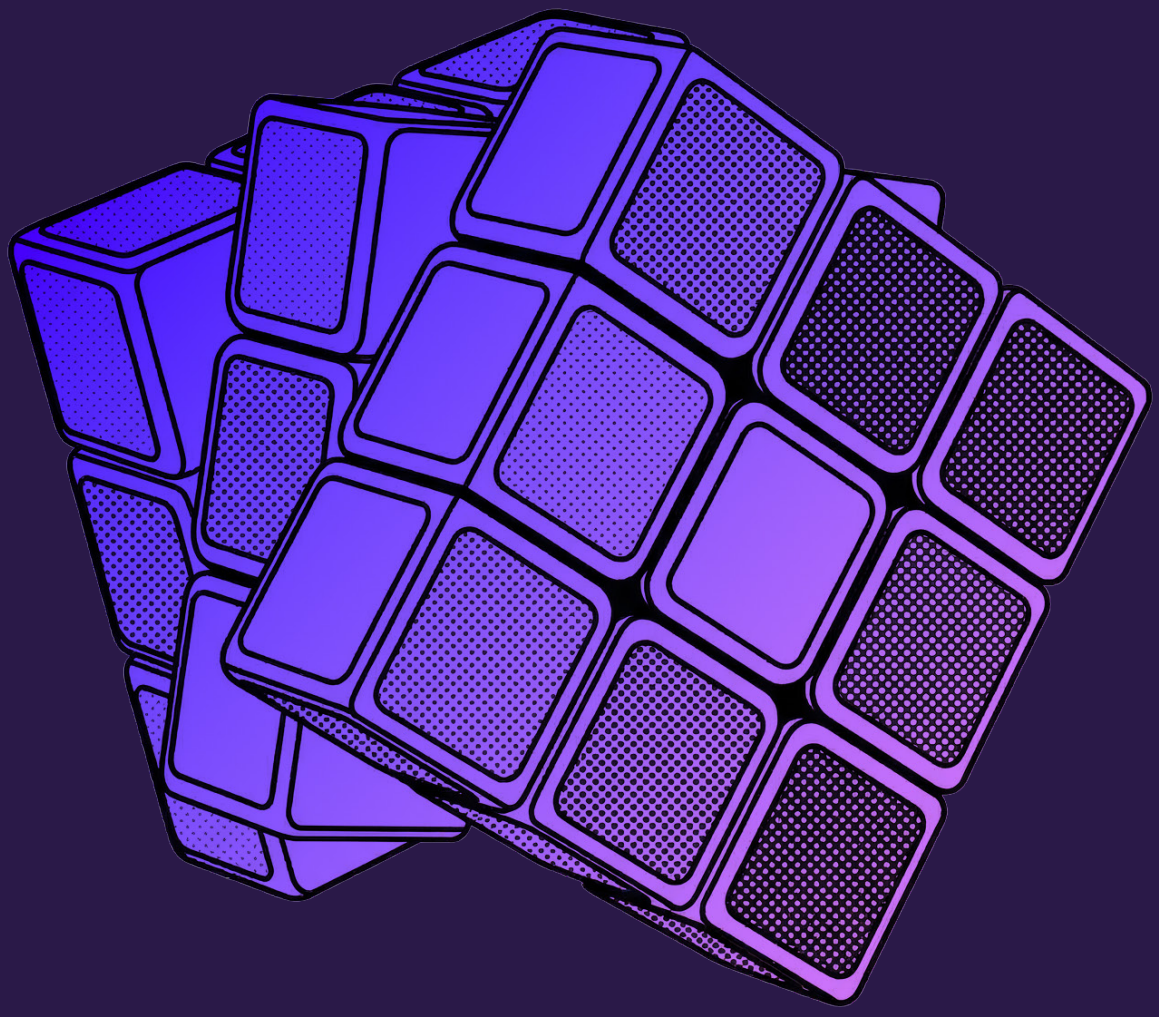

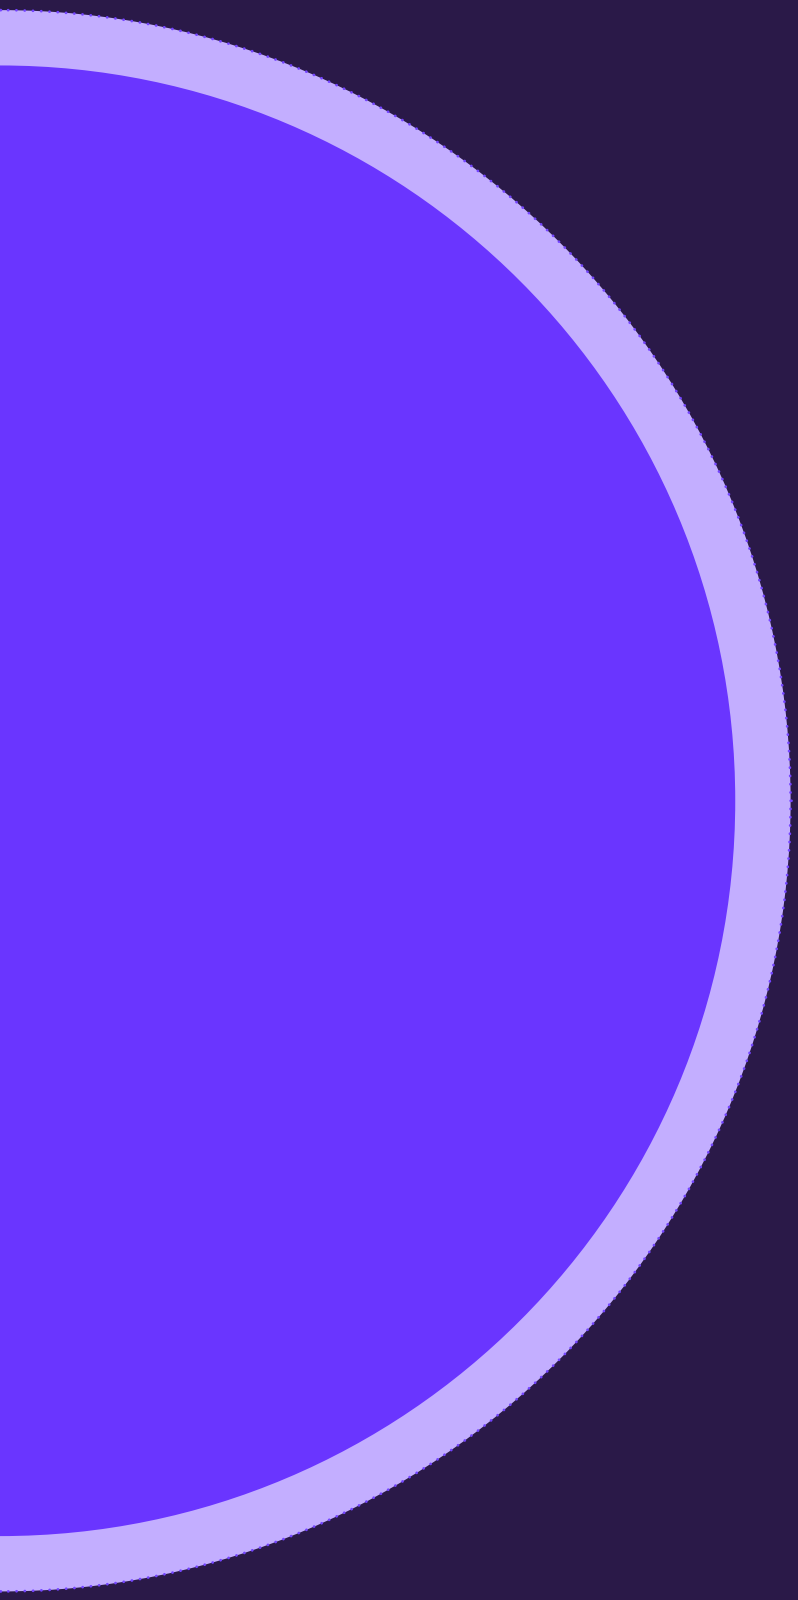

The average B2B company's market changes every month

7%

Here's how it works for 10,000 accounts in a 12-month period

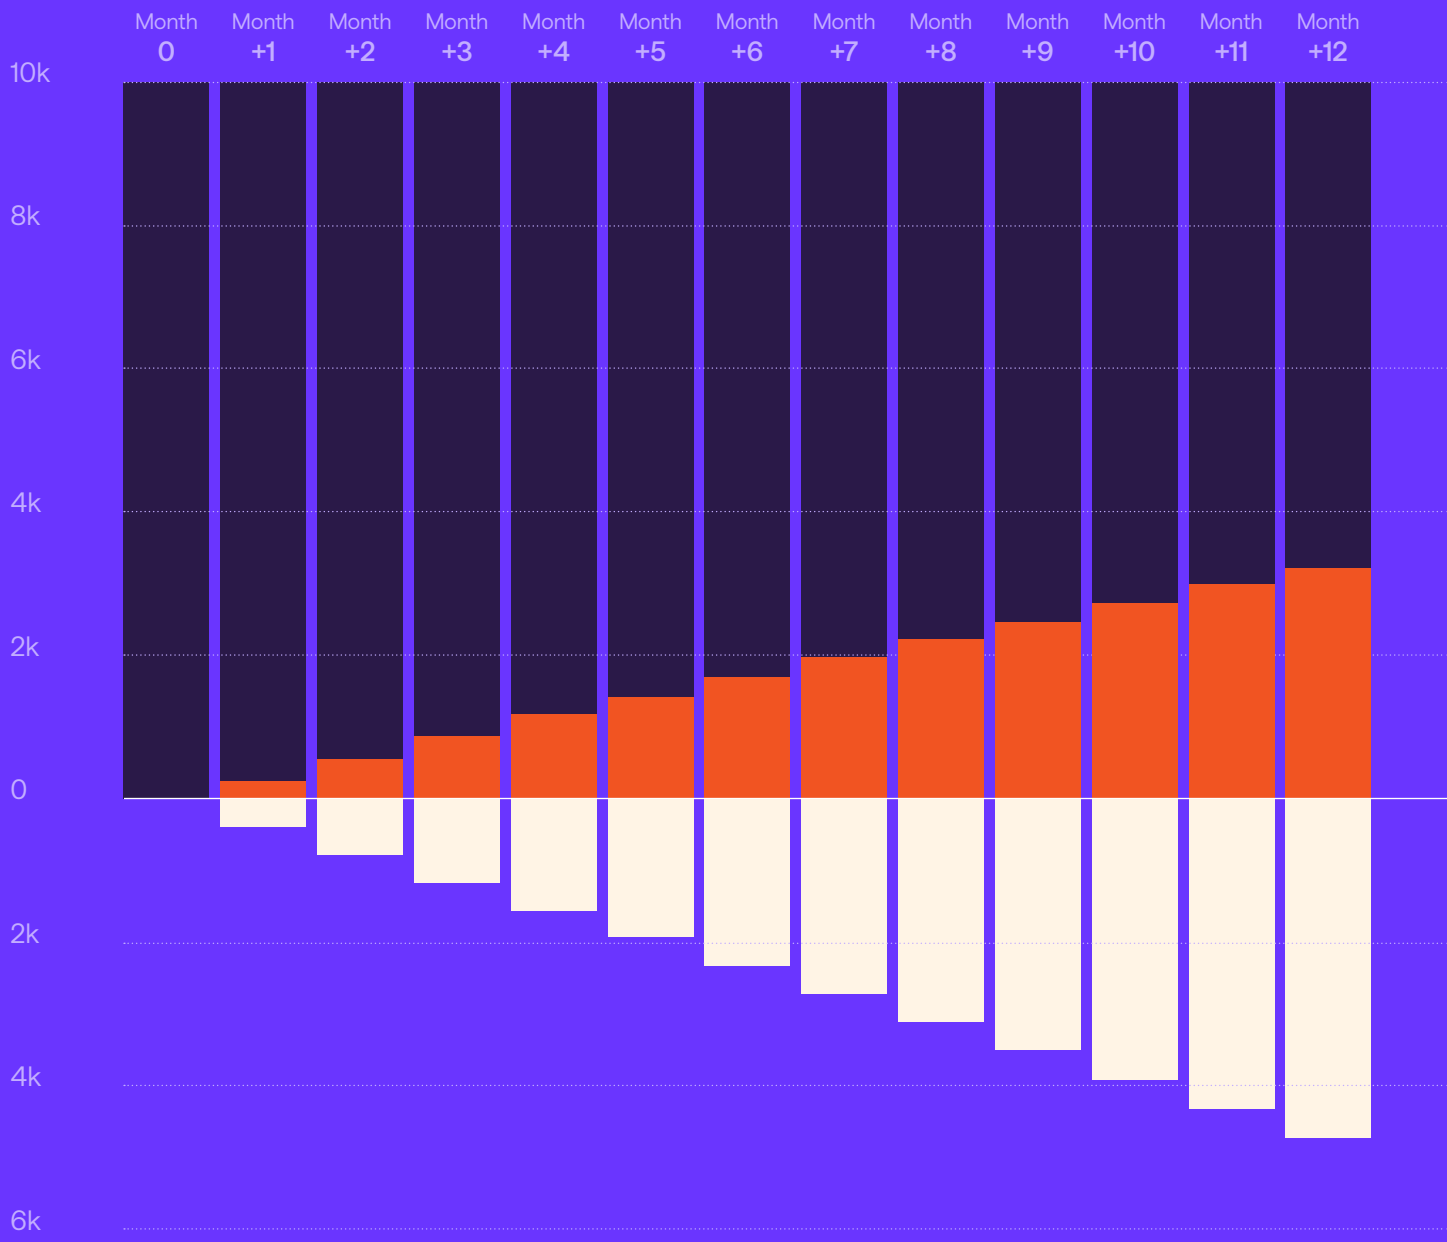

Accounts at Month 0	Still accurate at Month +12	False targets at Month +12	Invisible market at Month +12
10,000	6,767	3,233	4,714

A final and critical point: your market is not static. New accounts are constantly qualifying into your market.

A company hires their 10th sales rep, or opens an office in your region, or launches an inbound motion. Existing accounts become unqualified just as often: they get acquired, pivot, fire their 10th sales rep, shut an office in your region...

GoodFit tracked over 500,000 unique accounts across our companies' defined markets.
On average, 7% of accounts in a company's addressable market change every single month. Within a year, the majority of your market looks different from when you first mapped it.

The rate varies depending on how your market is defined. Companies with stable firmographic criteria (geography, industry, size) see 2-4% monthly movement. Companies whose market depends on dynamic signals — technology adoption, team composition, hiring patterns — see significantly more. But no market is static.

If you treated market mapping as a one-off exercise six months ago, a meaningful portion of that market has already shifted underneath you.

Market mapping is not a project.

It is an ongoing operation.

This report has described three problems. Your ICP is probably built on proxies that don't match who actually buys your product. The accounts that match your actual ICP are invisible because they can't be found through traditional data. And whatever view of your market you do have is decaying every month.

Solving this requires three things:

1. **Define your ICP using data, not assumptions.** Pull a list of your best customers, enrich them with a deep set of company-level attributes, and let the data tell you what your actual ICP is. Test whatever proxies you're currently using: how many of your existing customers actually match? What criteria are you missing? What criteria do you think matter but don't? Everything downstream depends on getting this right.

2. **Find every account that matches your criteria — including the ones that don't exist in traditional databases.** If your ICP can be captured through better filters (team composition, technology stack, web traffic), use them. If your criteria requires answering a question that no structured data field can capture ("does this company process personal data?", "does this company offer integrations?"), build a classification model that answers the question at scale, with human-level accuracy. Either way, the output should be a comprehensive, qualified dataset of every account in your market — not whatever ended up in your CRM through legacy prospecting.

3. **Keep your market current.** When 7% of your market changes every month, mapping is not a one-off exercise. Your market definition should be applied continuously: new accounts qualifying in, stale accounts flagged, contacts refreshed. Your CRM should reflect reality rather than a snapshot from six months ago.

If the numbers in this report gave you pause, there's an easy way to find out what your own market actually looks like.

GoodFit can build you a market report that shows your true market size based on your actual ICP, how many of those accounts are already in your CRM (and how many you're missing), the estimated revenue value of the accounts that are invisible to your team, and how many accounts you have that don't match your ICP and are consuming rep time and marketing budget.

It's free, and takes minimal work from your side. You can sign up at the link here.

GoodFit helps B2B revenue teams define, source, and prioritise every account in their market.

We validate what a qualified account looks like from your data, source the accounts that data providers can't surface, classify them with human-level accuracy, and keep that market current as it changes. So your GTM is built on your complete, qualified market — not whatever survived a set of filters.

GoodFit