siliconmint

AI-Powered Last Mile Delivery Platform for Urban Logistics

Smarter dispatching through AI-driven route optimization and real-time learning

We developed a full-featured last-mile delivery platform for a European logistics company specializing in bicycle-based deliveries. The system serves major brands like McDonald’s, Burger King, and local postal providers, supporting thousands of daily deliveries across several cities. Our solution combines complex B2B integrations, dynamic constraint-based dispatching, and a self-learning AI engine to optimize routes with minimal human intervention.

siliconmint

Business Challenge

- To scale bike-based delivery in **dense European cities**, the client needed to grow operations without overwhelming dispatch managers.
- Deliveries ranged from fast food (e.g., McDonald's, Burger King) to refrigerated medical packages — each with strict handling requirements.
- Manual route planning could only optimize **~20%** of deliveries
- **Dispatchers struggled to manage:**
  - Diverse delivery conditions
  - Varying courier capabilities
  - Tight service windows

siliconmint

Technical Challenge

Route planning had to consider not only geographic distribution but also:

Strict delivery windows:

two-hour slots, morning/evening windows, or ASAP same-day orders.

Item-specific constraints:

refrigeration, fragility, weight, temperature control, and urgency.

Vehicle compatibility:

standard bicycles, e-bikes, and electric tricycles — each with different terrain and load capacity.

Terrain complexity:

routes had to balance uphill/downhill/flat paths to prevent rider fatigue.

Dynamic order flow:

new deliveries added mid-route required real-time recalculations.

Manager knowledge:

only humans could initially detect impractical combinations (e.g., stacking three uphill stops for a courier with a heavy load and no electric assist).

This complexity outpaced basic algorithms, requiring a smarter system that could learn from real operations and continuously adapt.

siliconmint

Model definitions

MODEL NAME	ARCHITECTURE	PROVIDER	PURPOSE	INPUT	OUTPUT	USAGE CONTEXT
RouteAccept-LGBM	LightGBM (Gradient Boosting Decision Trees)	Azure ML	Predicts whether a route will be accepted, the reason for correction, and potential SLA risk	Structured delivery route features (~60 per route)	Classification label + score + ETA deviation estimate	Core dispatch logic, live route filtering
RouteScore-RNN	GRU-based recurrent neural network	Custom	Evaluates sequential logic and terrain consistency across the delivery path	Ordered stop list with metadata (20–30 steps)	Scoring value + flag for anomalies or fatigue risk	Supplementary scoring layer for complex routes
FeedbackParse-BERT	Transformer (BERT-style)	Azure ML / Hugging Face	Extracts structured issues from unstructured text comments	Text comments (~512 tokens) from clients, couriers, or dispatchers	Classified labels (e.g. “access problem”, “equipment missing”)	Post-delivery feedback analysis, route quality loop
ExplainRoute-GPT	GPT-4 (Generative Transformer)	Azure OpenAI	Generates a natural-language explanation of how and why a route was constructed	Route metadata, constraints, vehicle & item context	Full sentence-level summary of routing logic	Dispatch UI assistant, AI transparency layer

siliconmint

Runtime characterictics

MODEL NAME	AVG INFERENCE TIME	RETRAINING FREQUENCY	CONTEXT VOLUME	EXPLAINABILITY	REAL-TIME COMPATIBLE
RouteAccept-LGBM	~300 ms	Weekly	Fixed vector (~60 features)	✅ Feature importance	✅ Yes
RouteScore-RNN	~400 ms	Monthly	Up to 30 delivery steps	❌ Internal only	✅ Yes
FeedbackParse-BERT	~800 ms	Monthly	~512 tokens	✅ Label classification	⚠️ Partially (batch preferred)
ExplainRoute-GPT	~1–1.2 sec	No retraining (prompt-based)	~2,000 tokens	✅ Full natural-language output	⚠️ UI use only (not for routing logic)

Tech Stack

Azure Machine Learning

React

Azure SQL Server

Google Maps API

.NET

Angular

siliconmint

AI Optimization Logic

The AI engine was built as a self-improving assistant for dispatch managers — capable of learning from real-world corrections and continuously improving route quality. The goal was not just to automate routing, but to embed decision-making patterns observed in experienced dispatchers.

Structured Input & Feature Extraction

Each proposed route was transformed into a structured format, breaking down the sequence of delivery points with all relevant constraints.

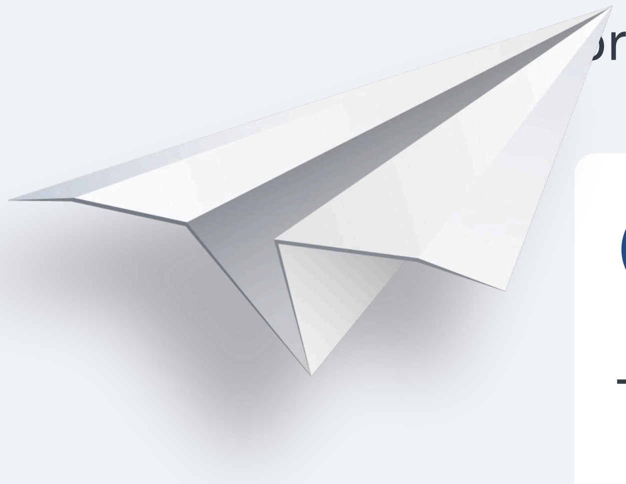

- Time sensitivity (e.g. ASAP, fixed slot)
- Item-specific rules (e.g. medication must be delivered first)
- Terrain and elevation complexity
- Vehicle capacity vs. load requirements
- Risk of courier fatigue based on route shape
- Compatibility between item and vehicle type (e.g. refrigeration needs)

These features were then aggregated to form a high-quality input set used for training machine learning models.

Learning from Human Corrections

We analyzed how dispatchers manually edited routes (reordering deliveries, reassigning couriers, removing overloads)
and identified recurring correction patterns.

Examples included:

- Medical deliveries placed too late in the sequence
- Too many consecutive uphill stops for non-electric bikes
- Mismatch between parcel size and courier equipment
- Overlapping time slots leading to potential SLA violations … etc.

This historical decision data allowed us to assign meaning to rejections and create labeled training sets that reflected both acceptance outcomes and the reason behind each correction.

LightGBM was chosen for its speed, accuracy, and explainability — enabling clear reasoning behind AI decisions.

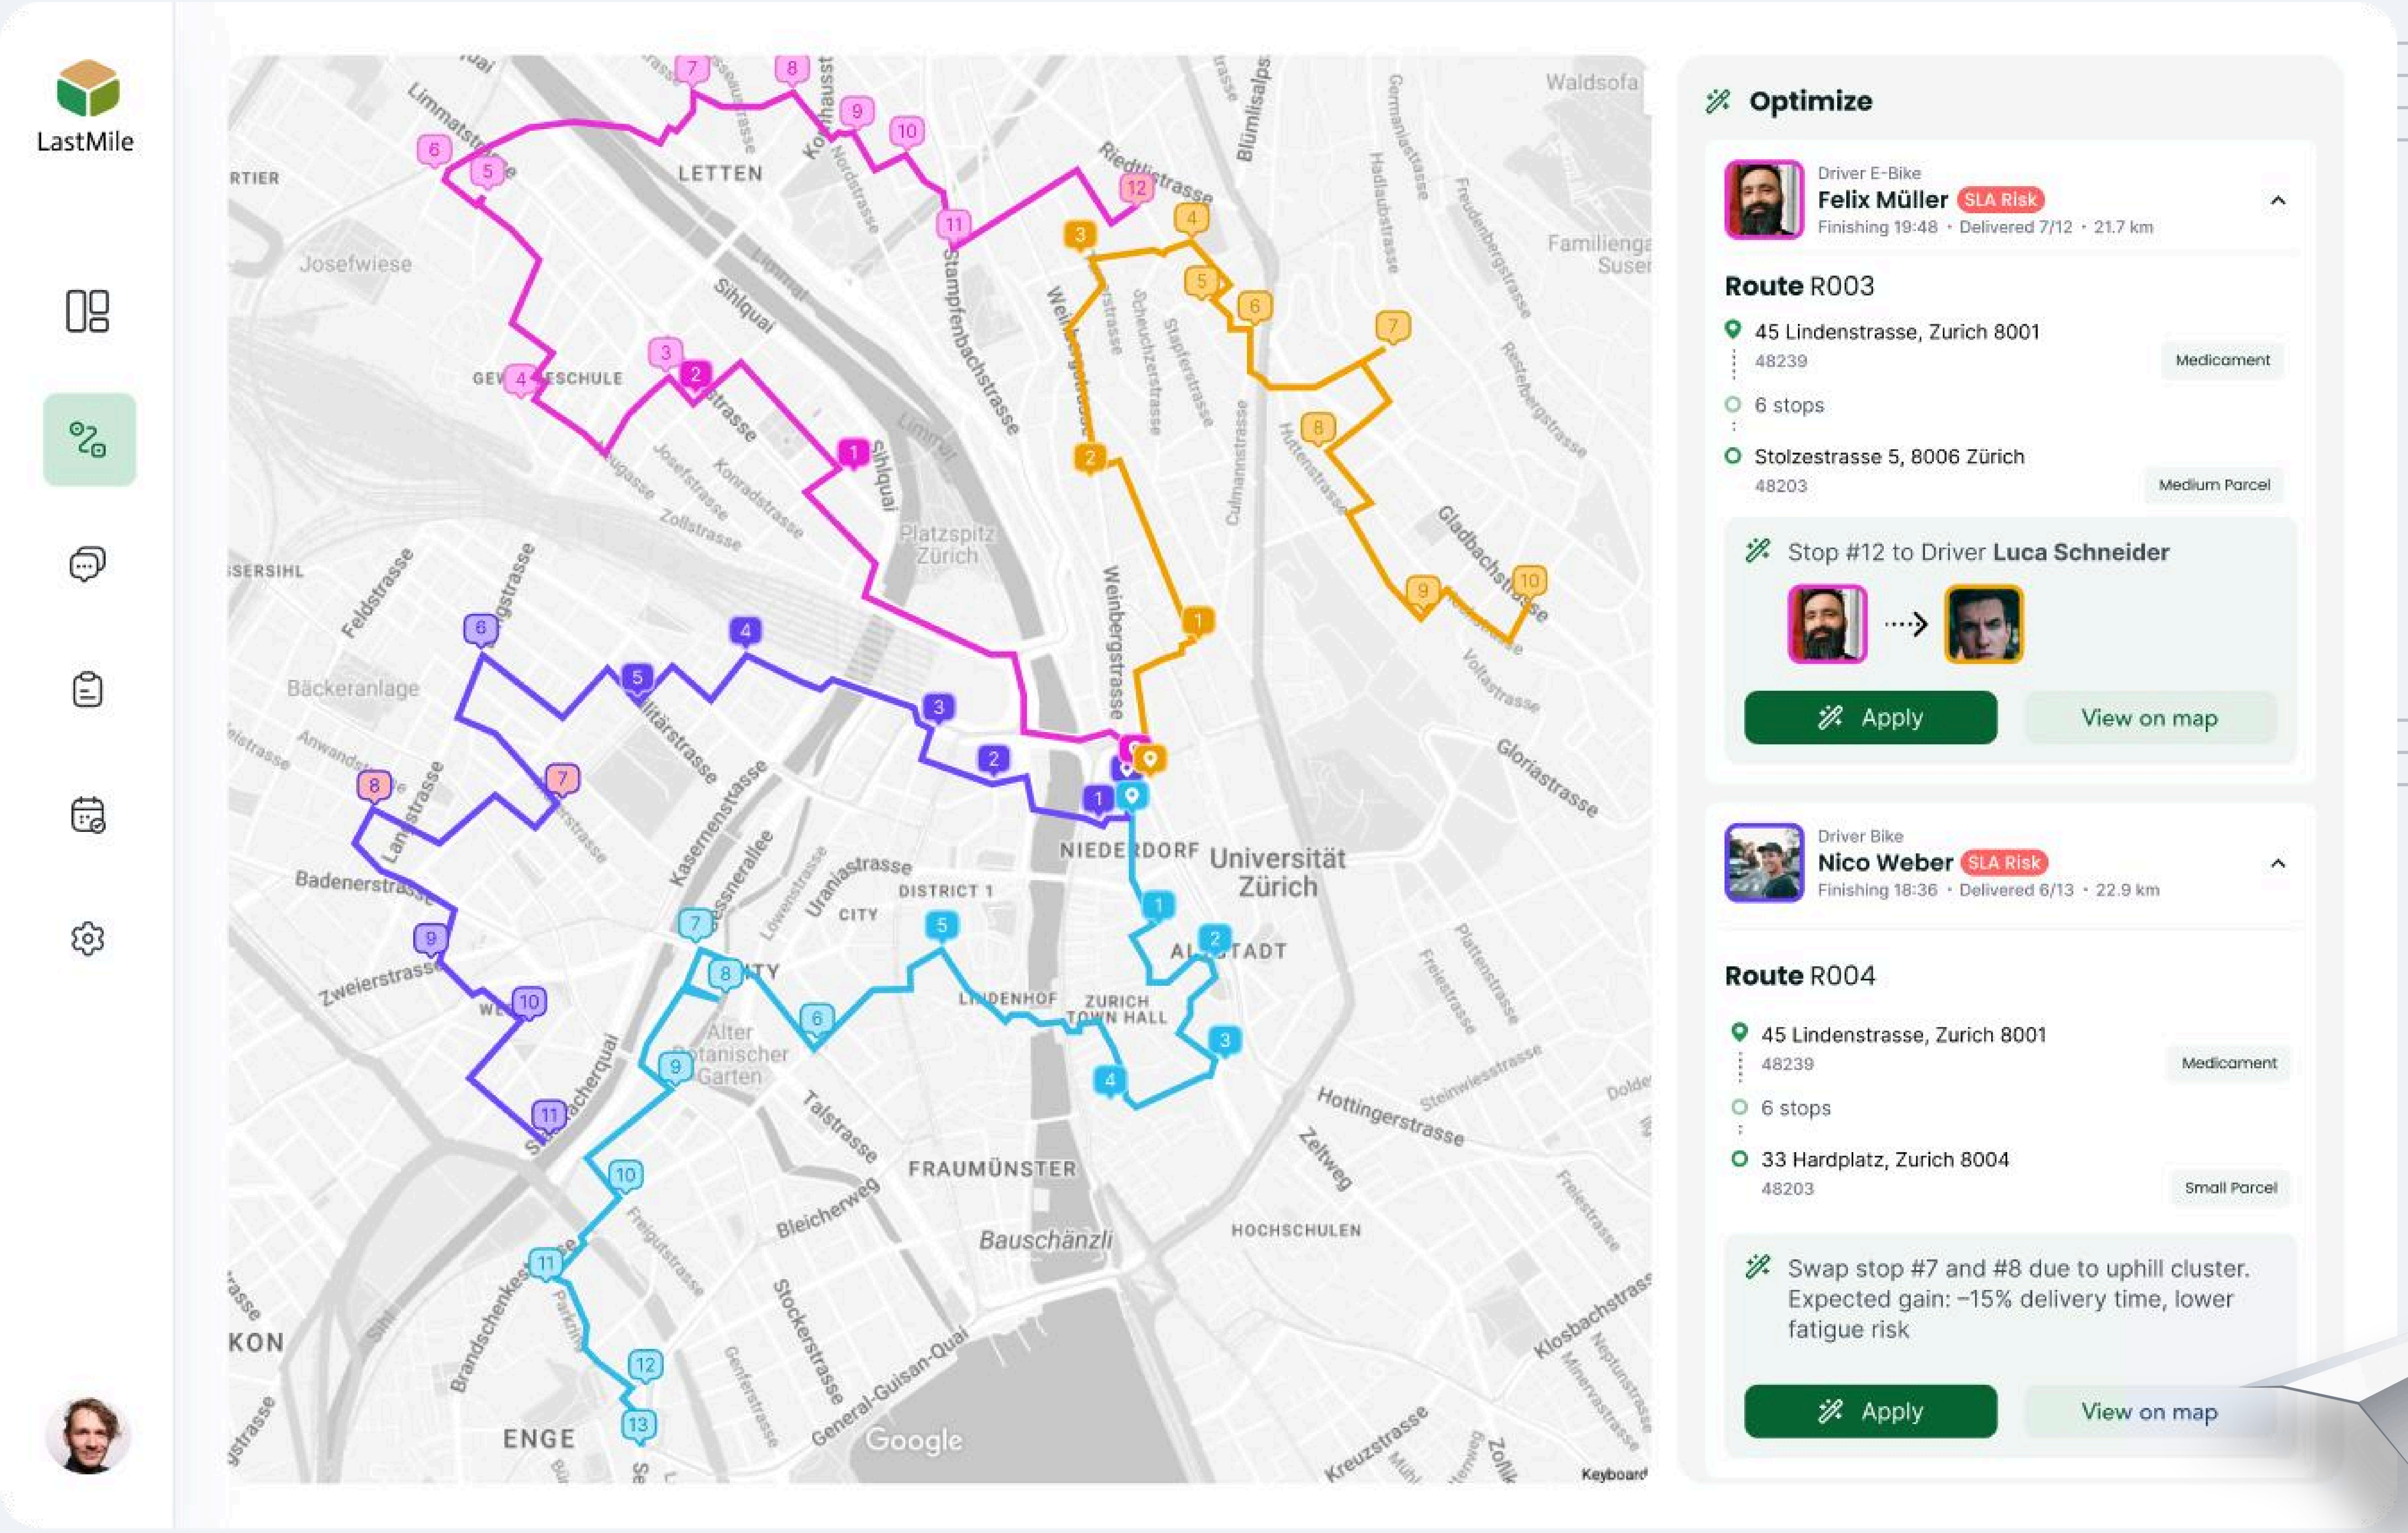
LastMile
Optimize
Driver E-Bike
Felix Müller
SLA Risk
Finishing 19:48 · Delivered 7/12 · 21.7 km
Route R003
45 Lindenstrasse, Zurich 8001
48239
Medicament
6 stops
Stolzestrasse 5, 8006 Zürich
48203
Medium Parcel
Stop #12 to Driver Luca Schneider
Apply
View on map
Driver Bike
Nico Weber
SLA Risk
Finishing 18:36 · Delivered 6/13 · 22.9 km
Route R004
45 Lindenstrasse, Zurich 8001
48239
Medicament
6 stops
33 Hardplatz, Zurich 8004
48203
Small Parcel
Swap stop #7 and #8 due to uphill cluster. Expected gain: −15% delivery time, lower fatigue risk
Apply
View on map
LETTEN
Josefwiese
NIEDERDORF
Universität Zürich
ALTSTADT
DISTRICT 1
LINDENHOF
ZURICH TOWN HALL
FRAUMÜNSTER
Bauschänzli
HOCHSCHULEN
ENGE
Google

Model Training & Adaptation

To train our AI models, we used Azure Machine Learning Pipelines — which allowed us to automate the entire flow from raw data extraction to weekly retraining and model evaluation.

We selected LightGBM (Light Gradient Boosting Machine) as the core algorithm for all major tasks. LightGBM is a fast, interpretable, and highly efficient model architecture specifically designed for structured data, making it ideal for learning from delivery routes, vehicle metadata, and historical correction patterns.

The models predicted:

- Whether a route would be accepted or edited
- The likely reason for rejection
- Estimated risk of SLA failure

LightGBM's native support for feature importance also allowed us to visualize the key factors influencing each decision, making AI outputs transparent and easy for dispatchers to trust.

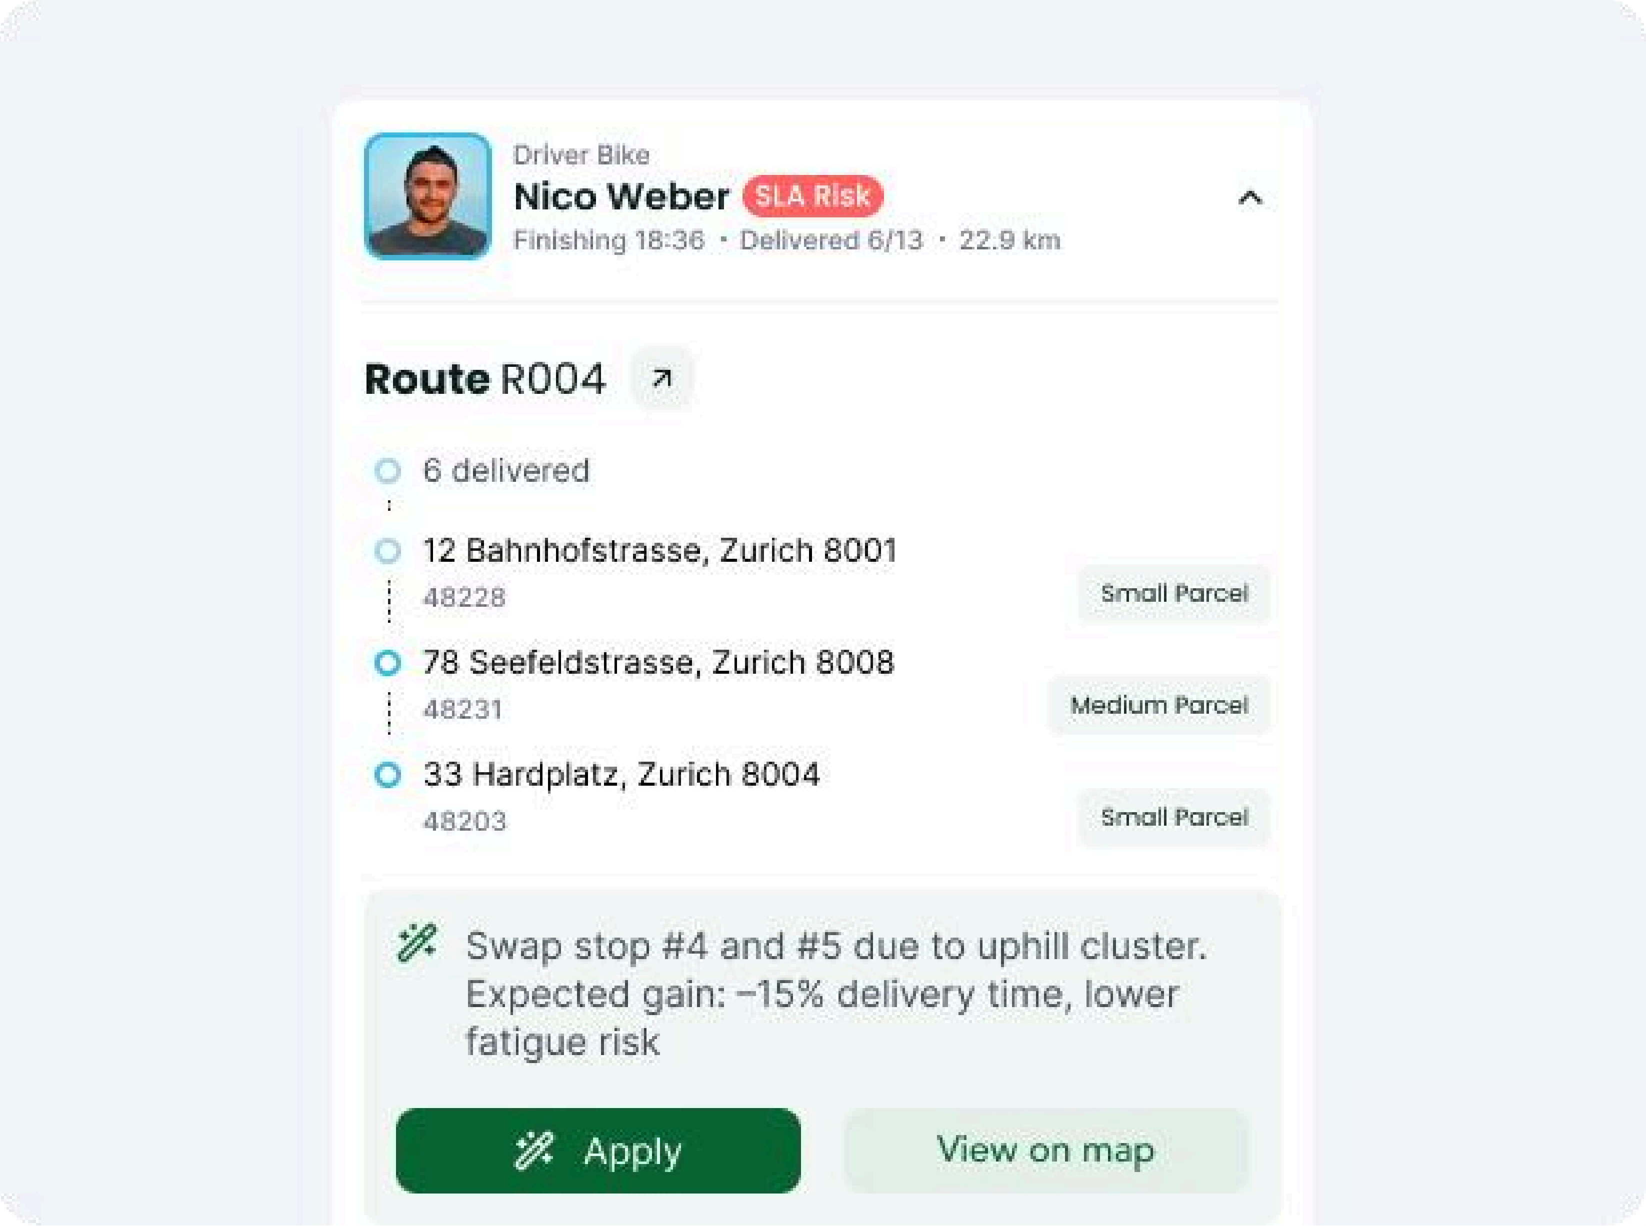

By the second month of use, more than 90% of routes were approved without manual adjustments.

Over time, the assistant became highly reliable — surfacing route suggestions that matched human preferences with minimal need for correction.

Neural Enhancements

To support richer understanding and human-AI collaboration, we added several lightweight neural components:

Route Scoring Assistant,

based on GRU architecture, analyzed routes as delivery sequences to capture hidden inconsistencies missed by flat models.

Feedback Interpretation Model,

(using BERT-style transformer) processed unstructured comments from dispatchers, couriers, and clients to extract insights from real-world events.

Generative Route Explainer,

powered by Azure OpenAI, produced natural language summaries explaining why a route was built in a certain way — improving transparency and trust in AI suggestions.

siliconmint

Our Solution

We delivered a robust enterprise delivery platform with:

Self-learning route engine

Powered by Azure Machine Learning, trained on real dispatch data and dynamic delivery constraints.

Terrain-aware route modeling

Built on Google Maps with elevation and traffic data for precise path planning.

Courier-aware assignment logic

Matches orders to bikes, e-bikes, or tricycles based on capacity, terrain, and load type.

Rich delivery metadata

Supports time windows, item fragility, refrigeration needs, and delivery priorities.

Seamless B2B integrations

Connected to food chains, postal providers, and third-party platforms.

Role-based access control

Tailored workflows for dispatchers, riders, support staff, and clients.

Unified backend on Azure SQL Server managing:

- Users
- Roles
- Permissions
- Delivery data
- Vehicle profiles
- Route histories
- Log storage
- Auditing
- Performance analytics
- SLA compliance
- Real-time status tracking

siliconmint

LastMile

Deliveries (234)

Today Last 7 days Last 30 days

Add order Import Orders

Search

Show only suggestions Hide delivered Display Columns Filter

Order ID	Date	Expected time	Route	Type	Client	Address	Driver	Status	Actions
48239	May 30, 2023	10:30 AM	R54321	Regular	MacDonald's	303 Elm Lane, Zürich	Jordan Parker	In progress	
75914	May 30, 2023	2:15 PM	R12345	Regular	Burger King	101 Birch Street, Zürich	Samantha Carter	Delivered	
23678	May 30, 2023	6:45 PM	R98765	Medicament	Healing Pharma...	404 Maple Road, Zürich	Avery Wilson	Scheduled	
91482	May 30, 2023	6:45 PM	R98765	Medicament	Healing Pharma...	404 Maple Road, Zürich	Avery Wilson	Scheduled	
58320	May 30, 2023	6:45 PM	R98765	Medicament	Healing Pharma...	404 Maple Road, Zürich	Avery Wilson	Scheduled	
73529	May 30, 2023	10:30 AM	R54321	Regular	MacDonald's	303 Elm Lane, Zürich	Jor		
86241	May 30, 2023	10:30 AM	R54321	Regular	MacDonald's	303 Elm Lane, Zürich	Jor		
64731	May 30, 2023	6:45 PM	R98765	Medicament	Healing Pharma...	404 Maple Road, Zürich	Ave		
89054	May 30, 2023	6:45 PM	R98765	Medicament	Healing Pharma...	404 Maple Road, Zürich	Ave		
32167	May 30, 2023	6:45 PM	R98765	Medicament	Healing Pharma...	404 Maple Road, Zürich	Ave		
39075	May 30, 2023	10:30 AM	R54321	Regular	MacDonald's	303 Elm Lane, Zürich	Jordan Parker	Scheduled SLA Risk	
47890	May 30, 2023	6:45 PM	R98765	Medicament	Healing Pharma...	404 Maple Road, Zürich	Avery Wilson	Scheduled	
15973	May 30, 2023	6:45 PM	R98765	Medicament	Healing Pharma...	404 Maple Road, Zürich	Avery Wilson	Scheduled	

Route R004 SLA Risk

Driver
Darrell Steward
Finishing 17:02 · Delivered 5/12 · 21.7 km

Swap stop #4 and #5 due to uphill cluster. Expected gain: −15% delivery time, lower fatigue risk

Apply View on map

siliconmint

Impact

98% AI route accuracy
after 2 months of learning from human feedback

85% reduction
in manual dispatch effort

35% drop
in courier fatigue incidents via smarter terrain-aware distribution

42% improvement
in SLA compliance across all delivery types

12+ cities and 5+ enterprise partners onboarded
through fast and scalable deployment

Zero rebuilds required
architecture scaled seamlessly with business growth