


# *AI-Powered* Offer Generation for the Automotive Market



# Table of Contents

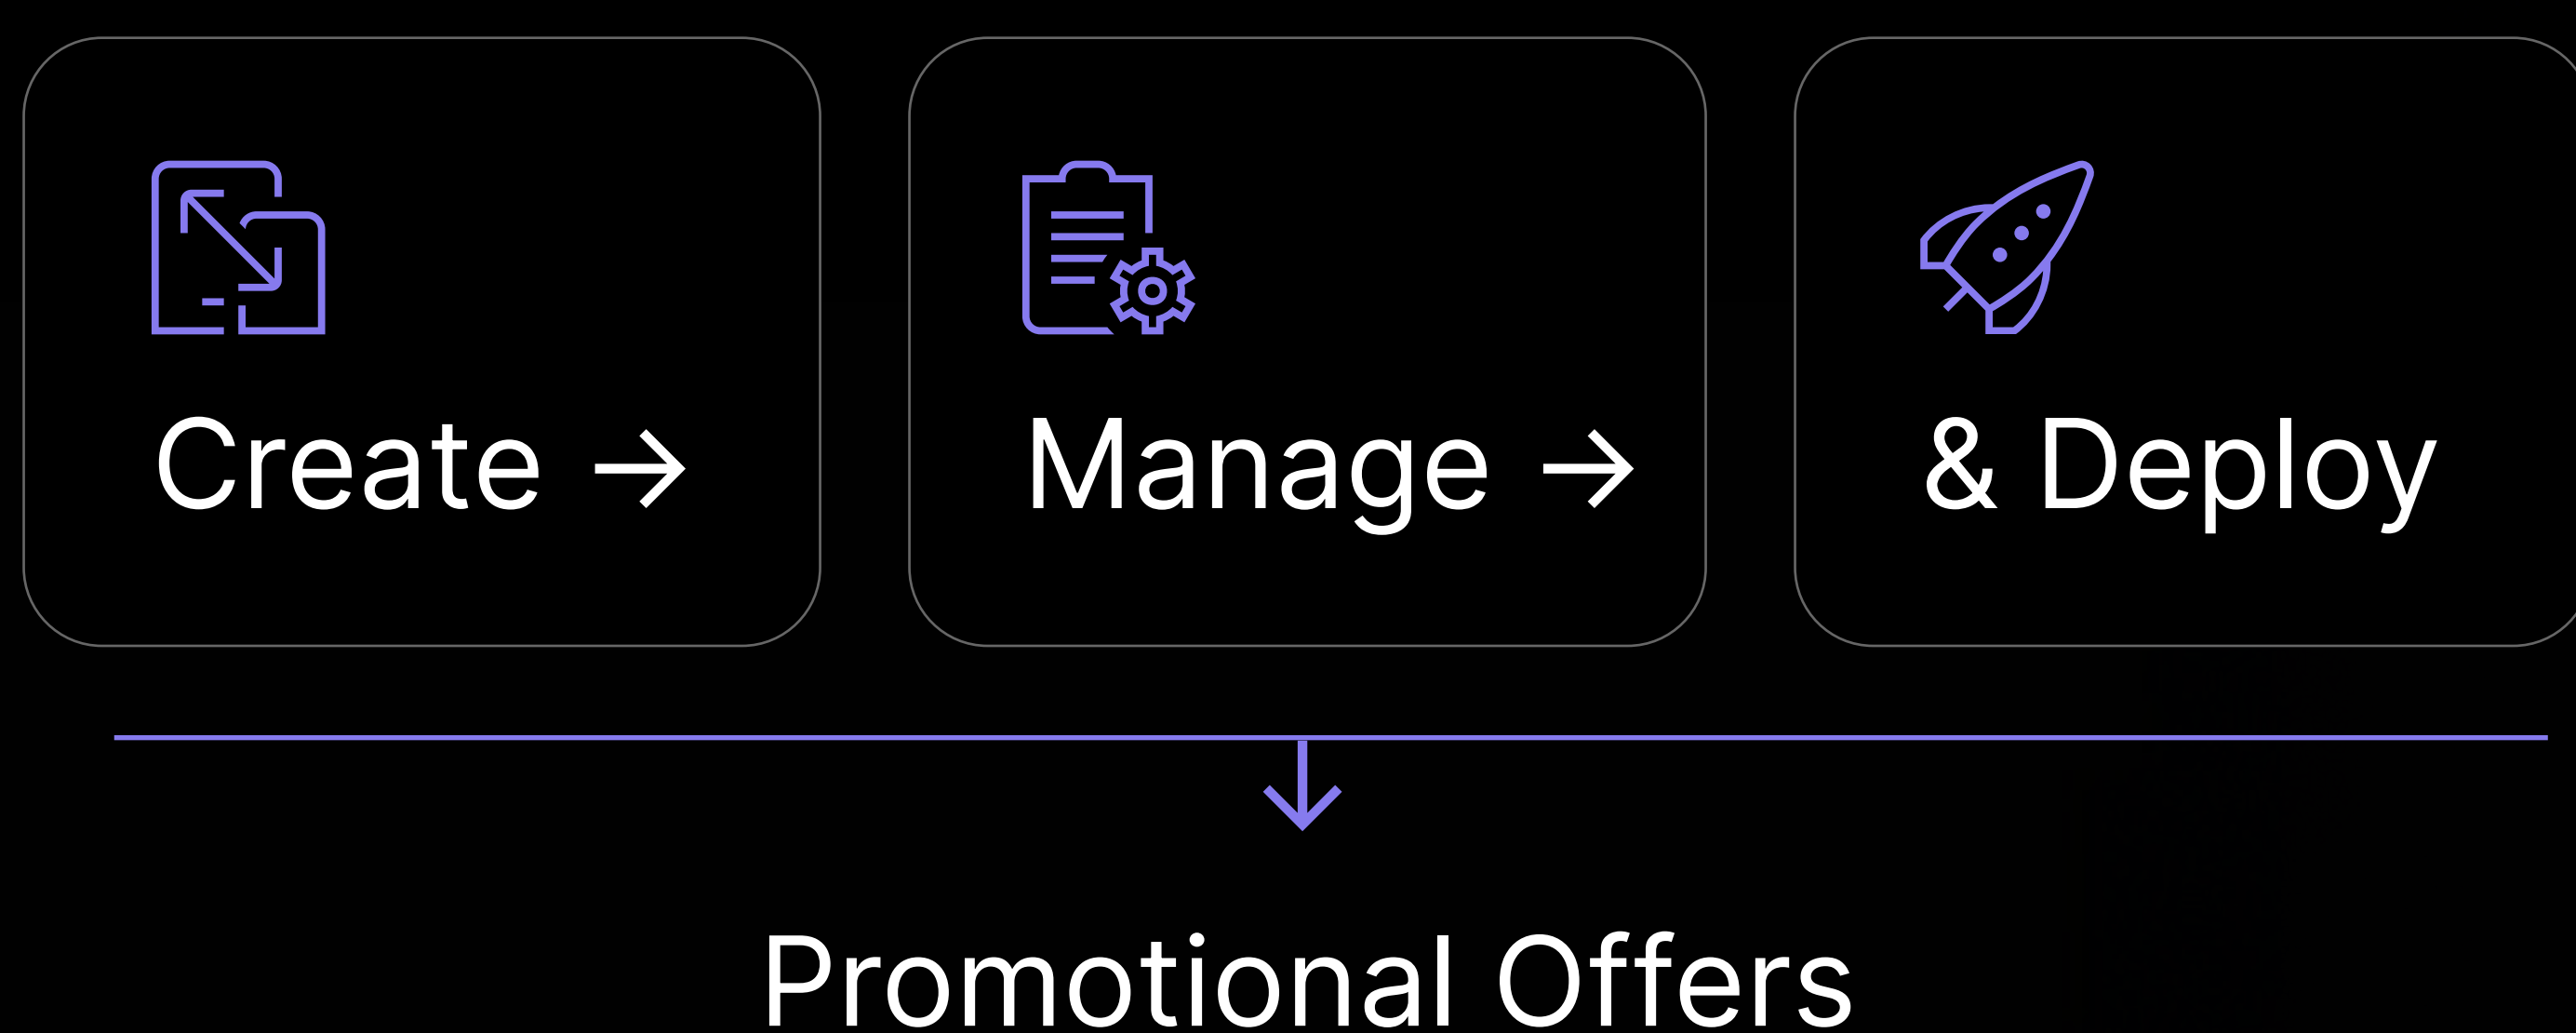
Executive Summary	02
Reimagining the Future of Offer Creation	03
Background	04
Proposed Workflow	06
Practical Integration & Continuous Optimization	07
Conclusion	08

# Executive *Summary*

Most automotive offers fail to drive action not because the value isn't there, but because the offer content doesn't fully engage the buyer. A headline that works for a December lease deal won't be the best fit for a spring tire change offer.

Today's customers expect personalization, but marketers are often reusing static templates – "Get a Deal on Oil Change" often becomes "Get a Deal on Ford Focus".

With thousands of promotional offers presented to customers nationwide daily, understanding what drives engagement is more important than ever. Launch Labs' marketing platform, Ignite, provides hundreds of automotive companies with tools to:



The platform also provides comprehensive performance analytics, which form the foundation for smarter, data-driven content recommendations.

This white paper discusses the applications and methodology of dynamic offer content generation, aiming to improve customer engagement, and make offer creation faster, more personalized, and grounded in what performs well.





*Reimagining*  
the Future of Offer  
Creation

Imagine a marketing platform where offers write themselves—tailored to the seasonal promotion, region, audience, and shopper profile. It's time to replace the guesswork in offer creation with data-backed innovation: **dynamic smart recommendations**.

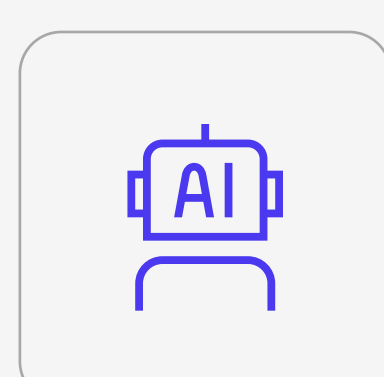
By using large language models and NLP techniques with high-converting offer data, Ignite can recommend content that reflects what has performed well previously. No more manually rewriting offer copies. No more wondering which headline to use.

## Benefits to Marketers:




### Higher-performing offers

tailored to specific audiences



### Faster content generation

for marketers



### Adaptation

to customer trends and seasonal patterns

# Background

## BACKGROUND

Ignite is Launch Labs' audience engagement platform designed to manage marketing and transform website visitors into **qualified sales leads**. In addition to managing promotional offers, it tracks offer engagement data including impressions (views), claims, conversations, appointments, and more. There is also associated metadata such as:



Audience Segment (New Vehicles, Used Vehicles, Service)



Offer type (APR, Discount, Trade-In)



Website type, run dates, offer images

By leveraging first-party data, Ignite allows businesses to collect valuable visitor information without intrusive forms of tracking, gaining insights into their audience while respecting privacy. This data includes shopper preferences, geolocation, and buying timeframes, which can be used to deliver personalized, relevant offer content. The platform integrates seamlessly with widely used CRMs, ensuring that visitor profiles are converted into actionable leads and sent directly to sales teams.

## What Are LLMs?

Large Language Models (LLMs) are AI systems trained on massive amounts of text to generate human-like language. Built on transformer neural networks with billions of parameters, they predict and assemble words into coherent responses.

## How it Works: AI-Powered Smart Recommendations


Large Language Models (LLMs) are AI systems trained on massive amounts of text to generate human-like language. Built on transformer neural networks with billions of parameters, they predict and assemble words into coherent responses.



### OFFER HEADLINE

The offer headline should be a short, attention grabbing statement. Here's some things that have worked well for other users:

Our annual sale is going on now an exclusive offer just for you we've got a great deal for you



### LONG HEADLINE

The long headline should give information the product or service that the offer actually applies to

Get a great deal on new chevrolet express extended cargo van!

## BACKGROUND

A key performance indicator for offer success is the **claim conversion rate**, defined as the percentage of impressions that convert to claims on a given offer. To identify high-performing offers, we can merge historical engagement data (impressions, claims, etc.) with offer content (offer headline, long headline, etc), and then filter the data by highest claim conversion rate, organized by audience segments. For a given offer, we get the following metric:

**Claim Conversion Rate =**  
**(Number of Claims/Number of Impressions) \* 100**

For example, if we have 100 claims for a given offer, and that offer has 900 impressions, we have a conversion rate of  $100/900 * 100 = 11.11\%$ .

This metric allows us to identify what drives user interaction. We can look at conversion rate on any level, whether it be conversion to conversations or appointments, but in this context, claim conversion rate— whether or not a customer clicks— can tell us a lot about the effectiveness of the content of the offer.

## Few-Shot Prompting

Once we have identified our top-performing offers, we can utilize few-shot prompting, where we feed 3-5 real examples into a large language model (e.g. gpt-4-turbo), which learns the tone, structure, and phrasing of a successful offer. Using an appropriate prompt, we ask it to generate new content in a similar style. Then, using selected inputs, audience segment and a selected offer text, we ask the AI to generate new offer recommendations as identified below:




# Proposed *Workflow*

## Semantic Search


While conversion rates help surface strong examples, semantic search adds another layer of precision by finding offers that are not just high-converting, but also contextually similar in meaning to the current offer being created. This ensures the AI learns from examples that reflect both the content and intent of the offer.

## Key Concepts:



### Sentence Embedding

A numeric vector representation of a sentence's meaning, used to compare textual similarity.



### Semantic Search

Method of finding similar content based on contextual meaning rather than just keyword matches. It uses sentence embeddings to compare and identify matches, often measured using cosine similarity.

Each new offer is vectorized using an embedding model (e.g. text-embedding-3-small) and then compared to existing offer vectors. This allows the system to find the most relevant examples, which is useful for providing more specific recommendations within a broad audience segment. For example, semantic search can help distinguish the difference between a suitable headline for an oil change and a tire rotation offer, even though both are Service offers.

## USE OF AI

AI tools (LLMs) were used to support drafting and ideas, with all outputs reviewed by humans for accuracy and quality.

# Proposed *Workflow*

- 01 Select audience segment.
- 02 Select an offer text from recent high-performing offers.
- 03 Generate smart recommendations from the top 35 most contextually similar high-converting offers in the past month, by the given audience segment.
- 04 Construct a prompt using these examples, with instructions and desired output.
- 05 Send prompt via an API key to a commercially licensed large language model API (e.g. gpt-4-turbo)
- 06 The model will output 3 options each for headline, long headline, call to action, and other fields you wish to generate recommendations for.
- 07 Output is formatted into options that can be displayed directly in Ignite's Offer Builder interface or desired marketing platform. Optionally, there could be an implementation for a user to select a theme, such as a holiday promotion, for a more specialized offer recommendation.

# Practical Integration & Continuous *Optimization*

# Practical Integration & *Continuous Optimization*

## User interface


The updated user interface for the client will be similar to the current Ignite Offer Builder but will include 3 dynamic recommendations each for headline, long headline, call to action, and other desired fields. These will be updated on a regular basis based on recent successful offers.

## Ongoing Data Refresh


To keep recommendations relevant, the backend system will periodically refresh the offer data from an internal database. This ensures the model is always using the most recent highest converting offers for its recommendations. Update frequency could be daily or weekly, depending on user needs and data volume.

## Prompt Engineering & Model Fine-Tuning


We use prompt engineering and parameter tuning to continually refine and improve output recommendations. This includes:



Experimenting with prompt structure, review prompt-output behavior, and monitor efficiency in a secure LLM testing environment



Adjusting model parameters such as temperature, which represents the model's "creative freedom," and example count for clarity



Exploring tools to organize, compare, and manage prompt versions at scale

Smart offer generation represents a new chapter in automotive marketing, combining AI and natural language processing to drive innovation in content personalization and optimization. Real performance data can help marketers confidently respond to shifting consumer behavior. This approach makes content generation fast, efficient, and relevant for marketers, ultimately leading to boosted customer engagement and increased sales.

**launch**  **labs**