生成即分配：大模型原生广告中的词元级拍卖与机制设计

Generation as Allocation: Token-Level Auctions and Mechanism Design for LLM Native Advertising

祁琦
中国人民大学高瓴人工智能学院
2026 年 8 月 29 日

全局混排：AI 自动学习页面级广告机制

问题不只是“广告插在哪里”，而是**广告与自然内容如何共同组成一个页面**。

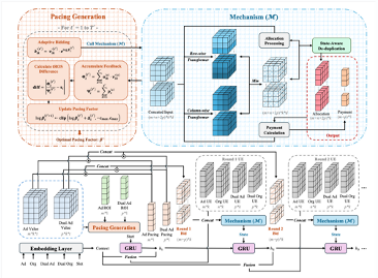

机制设计基础 [TCS23] · 去重与激励相容 [AAAI25a, Oral]
上下文外部性 [WWW25] · 自动出价 [KDD26]

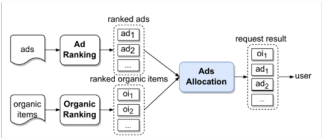

- 统一候选池：
  广告 + 自然内容
- 目标：页面级长期价值
- 约束：激励、预算与 ROI

AI 的作用：在既有页面范式内，自动搜索更复杂的排序与拍卖规则。

联合拍卖：AI 为共享商业意图设计新市场

同一次转化背后，店铺与品牌可能拥有**可共享、但并不相同的商业意图**。

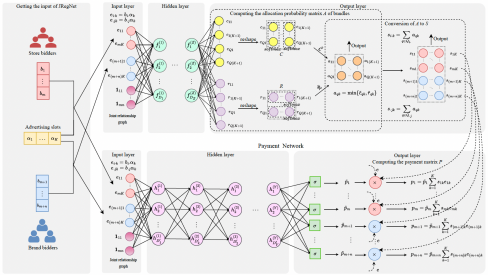

JRNet [KDD24 Oral] · BundleNet [ICML25]
HRegNet [SIGIR25] · JTransNet [KDD26 Oral]

- **品牌**推广某类商品
- **店铺**争取具体转化
- **平台**联合分配、耦合支付

AI 的作用：学习可部署、满足 DSIC/IR 的组合拍卖，扩张市场设计空间。

GenAuction：AI 直接生成页面级分配

当价值取决于页面上下文，机制可以从选择 winner 走向**生成整个页面**。

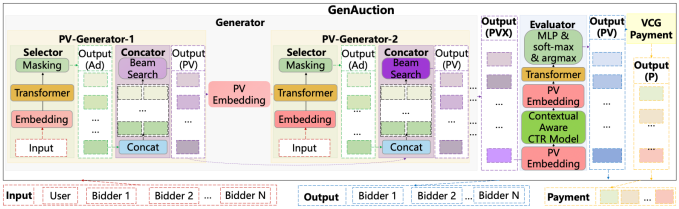

GenAuction [AAAI’25]

- 候选广告与页面上下文共同输入 Generator
- Evaluator 估计福利、约束与用户体验
- 以 CMDP 处理品牌冲突、预算等多目标约束

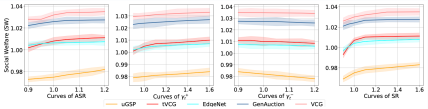

生成策略在多组约束曲线下保持稳定

AI 的作用：把机制设计从规则驱动拍卖推进到生成式机制设计。

MARSHALL MCLUHAN 1964

“The medium is the message.”

媒介不只是承载信息，
它也塑造信息如何被组织。

广告机制首先继承媒介的结构。

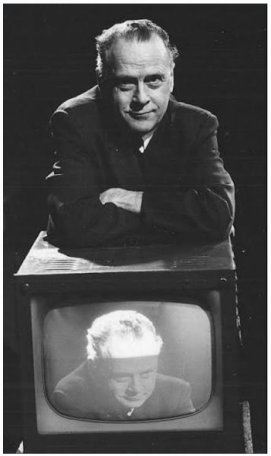

Marshall McLuhan

Token-Level Advertising

arXiv:2608.27382v1

Hanbing Liu
RUC

Bowei Zhang
RUC

Changyuan Yu
Baidu

Yinyu Ye
Stanford

Qi Qi
RUC

传统互联网广告围绕槽位生长

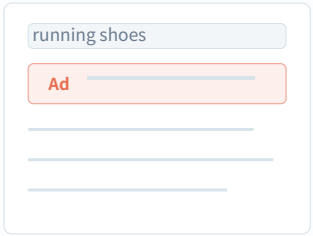

Sponsored search

排名列表中的固定位置

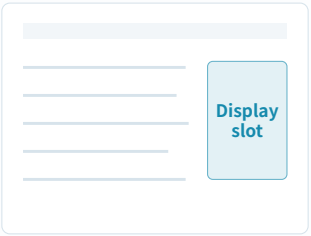

Display advertising

页面中预留的展示区域

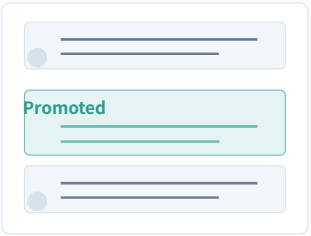

Promoted feed

信息流中的预设位置

媒介先组织内容 → 平台定义库存 → 拍卖分配槽位

三十年广告技术，都从固定广告位出发

2025 年全国广告业务收入 **2.05 万亿元**；互联网广告发布收入 **1.25 万亿元**

1990s

1994 · Banner
HotWired

1996 · Ad Server
DoubleClick

广告位开始可服务、可追踪

2000s

2000 · Search Ads
Google AdWords

2005 · Exchange

意图匹配与库存聚合

2010s

2010 · DSP / SSP / DMP

2012 · Programmatic
Autobidding

2014 · Video / Feed

生成式入口

2020s · AI answers

2026 · LLM Ads

媒介开始直接生成内容

界面不断变化，但核心抽象始终是：先定义广告位，再竞价分配。

现在，信息入口开始直接生成答案

百度 AI 智能回答

在搜索结果中直接生成综合回答

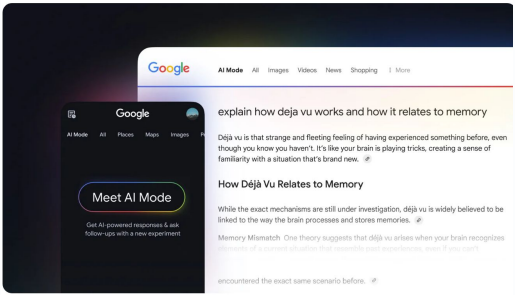

Google AI Search

在搜索入口直接合成回答

信息入口从“排序已有内容”变为**“围绕意图动态生成回答”**。

最直接的方案：把广告卡片挂在答案周围

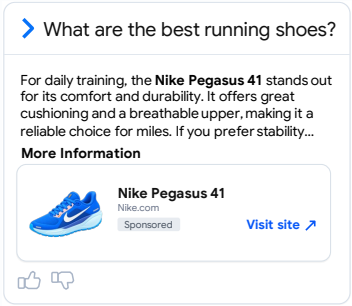

回答下方的 Sponsored Card

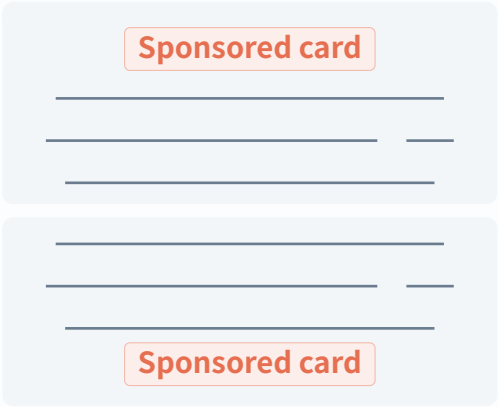

回答先生成，广告库存仍然独立存在。

但“外挂卡片”还没有形成规模化商业通道

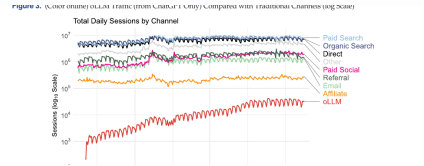

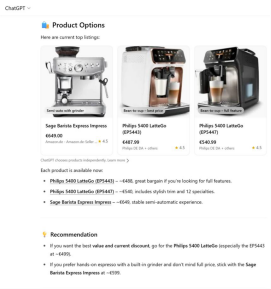

$< 0.2\%$

样本全部电商访问中，来自 ChatGPT referrals 的占比仅有

$\approx 1/200$

Google organic search 的访问量

实证数据显示，**从 ChatGPT 回答上的卡片跳转到商家**的点击率显著低于传统广告形态。

真正原生的广告机会，应当在生成中形成

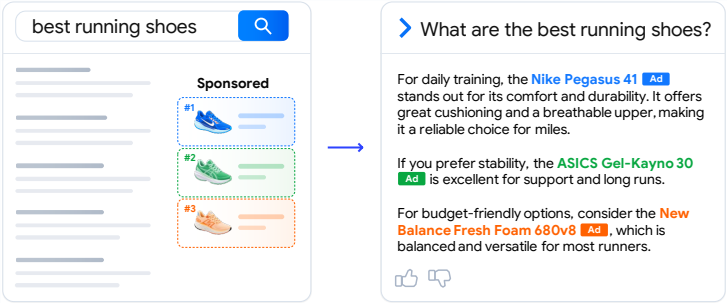

Generation does more than fill an opportunity—it can help create one.

内容生成 = 机会形成 + 资源分配

广告机会正在向生成过程内部移动

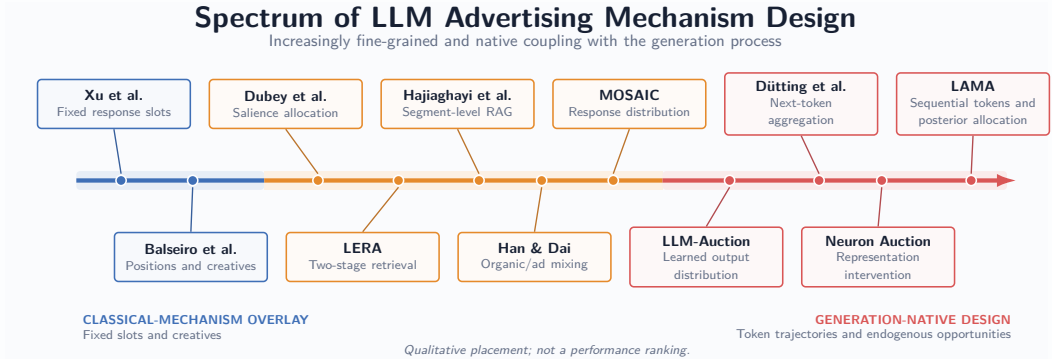

从固定槽位与完整回复，走向**位置、分布、隐表示与生成轨迹**。

代表性机制

Xu et al.

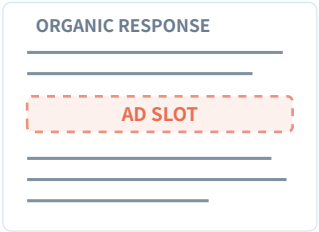

先生成自然文本

再将创意插入固定句间位置

Dubey et al.

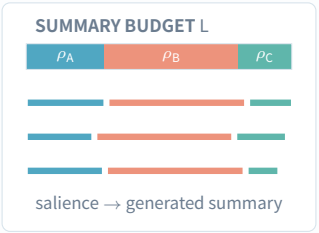

竞价先决定显著度

LLM 按**各广告主相对篇幅**生成

Dütting et al.

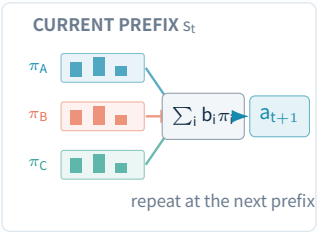

每一步聚合分布偏好

独立地决定当前 token

分配对象逐步细化：**slot** → **length share** → **next-token distribution**

极致原生：广告主参与每一个词元决策

For daily training, the Nike Pegasus 41 stands out;
for stability, consider the ASICS Gel–Kayano 30.

每个前缀

广告主报告局部信息，直接改变下一词元分布

每个词元

新文本同时更新“谁更可能成为最终赢家”的后验

结算

后验决定单赢家分配，并沿整条路径结算支付

广告主不是等待分配，而是全程参与生成本身。

回复树

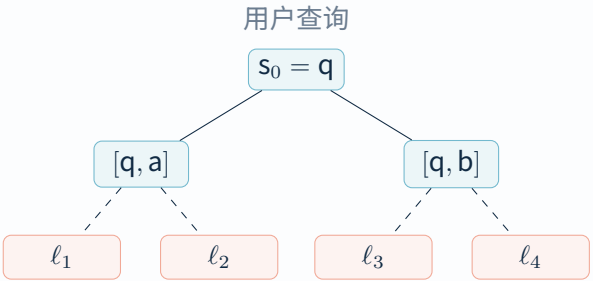

两个量随路径变化

- 回复自然度
- 广告主终端价值 $r_i(\ell)$

早期一个词元，可能改变后面所有广告机会。

模型：生成内容的展示价值

生成

$$
\begin{aligned}
s_{t+1} &= [s_t, a_t], \\
a_t &\sim \pi(\cdot \mid s_t), \\
\ell(\tau) &= s_T.
\end{aligned}
$$

分配与价值

$$
\begin{aligned}
r_i : \mathcal{S}_T &\to \mathbb{R}_{\geq 0}, \\
\lambda(\ell) &\in \Delta(\mathcal{N}), \\
\text{gross value}_i &= \lambda_i(\ell) r_i(\ell).
\end{aligned}
$$

π_{ref}：无广告时的用户中心参考策略

单赢家设置：机会由生成塑造，终局归于一名广告主

生成策略 π + **终局分配 λ**

目标：广告价值与自然生成之间的权衡

$$\mathrm{SW}_{\beta}(\pi, \lambda) = \mathbb{E}_{\tau \sim \pi}\left[\sum_{\mathrm{i}} \lambda_{\mathrm{i}}(\ell)\mathrm{r}_{\mathrm{i}}(\ell) - \beta \sum_{\mathrm{t}} \log \frac{\pi(\mathrm{a}_{\mathrm{t}} \mid \mathrm{s}_{\mathrm{t}})}{\pi_{\mathrm{ref}}(\mathrm{a}_{\mathrm{t}} \mid \mathrm{s}_{\mathrm{t}})}\right]$$

广告价值

- 谁获得终局机会
- 回复对该广告主是否有效

用户质量锚

- 轨迹级 KL 偏离成本
- β 越大，越接近自然回答

每个广告主先解一个软控制问题

$$V_i^*(s) = \max_{\pi} \mathbb{E}\left[r_i(\ell) - \beta \sum_{t \geq s} \log \frac{\pi(a_t \mid s_t)}{\pi_{\text{ref}}(a_t \mid s_t)} \right]$$

软贝尔曼递推方程

$$V_i^*(s) = \beta \log \sum_a \pi_{\text{ref}}(a \mid s) e^{V_i^*([s,a])/\beta}.$$

广告主最优策略

$$\pi_i^*(a \mid s) \propto \pi_{\text{ref}}(a \mid s) e^{V_i^*([s,a])/\beta}.$$

在线报告只需要“下一步的继续价值”

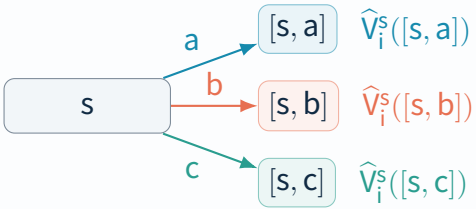

报告对象不是完整奖励函数，而是当前前缀下的子状态 value vector。

隐含了当前状态的局部最优策略

$$\widehat{\pi}_i(a \mid s) \propto \pi_{\text{ref}}(a \mid s) e^{\widehat{V}_i^s([s,a])/\beta}.$$

每个词元由一个“潜在广告主”引导

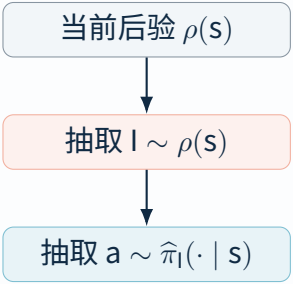

边际上等价于混合策略

$$x(a \mid s) = \sum_i \rho_i(s)\hat{\pi}_i(a \mid s).$$

I 不直接等于最终赢家；它只是当前词元背后的潜在来源。

生成出的词元也是关于最优广告主的证据

$$\rho_i([s,a]) = \frac{\rho_i(s)\widehat{\pi}_i(a \mid s)}{\sum_j \rho_j(s)\widehat{\pi}_j(a \mid s)}$$

生成前

$\rho(s)$

$\rightarrow$

观察到词元 a 后

$\rho([s,a])$

若广告主 i 更可能生成 a，它的后验权重就上升。

最终后验直接成为广告分配

$$\lambda(\mathsf{s}_T) = \rho(\mathsf{s}_T) \in \Delta(\mathcal{N})$$

分数不是对完整回复事后计算；它是生成中逐步积累的。

四步机制：报告、生成、更新、结算

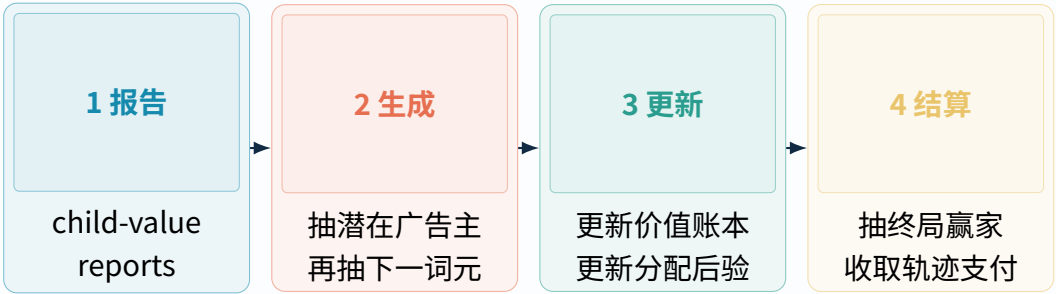

LAMA = Latent Advertiser Mixture Auction

“Latent Mixture” 得名于生成中的**贝叶斯推断**：逐词元更新潜在广告主后验。

机制运行流程

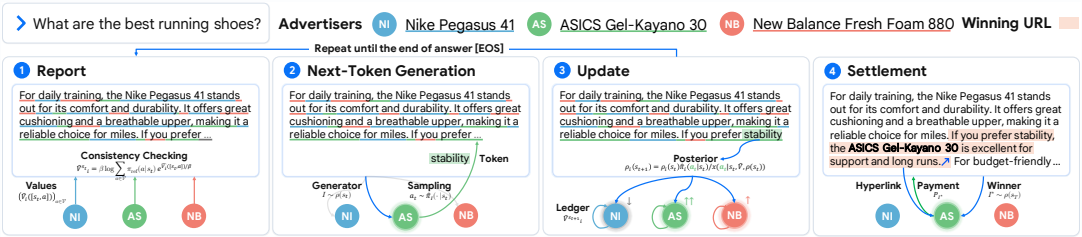

Consistency Checking? 第一步为什么不能让广告主在每个状态随心所欲地报告 child values?

贝尔曼一致性检验

$$\beta \log \sum_{a} \pi_{\text{ref}}(a \mid s) e^{\widehat{V}_i^s([s,a])/\beta} = \widehat{V}_i(s)$$

初始状态

- 可自由报告第一组 child values
- 平台记录根价值 $\widehat{V}_i(q)$

后续状态

- 新报告必须聚合回账本中的 parent value
- 只能改变子树内的分配，不能重写过去

Markov DSIC：有利前缀出现后仍不值得改报

定理

在贝尔曼一致的消息空间中，诚实报告

$$\left(V_i^*([s,a])\right)_{a\in\mathcal{V}}$$

对任意对手策略、任意可达前缀及任意可行 continuation deviation 都最优。

证明路线

1. Bellman 一致性使叶子状态的 $z_i(s_T) = e^{\widehat{V}_i(s_T)/\beta}$ 能沿生成树递推到根节点。
2. 终端“生成且分给 i”的概率是凸 log-sum-exp 势函数的梯度。
3. 势函数支付

$$P_i^{\mathrm{frac}}(s_T) = \rho_i(s_T)\widehat{V}_i(s_T) - \beta\log\sum_{j\in\mathcal{N}} e^{\widehat{V}_j(q)/\beta} + \beta\log\left(1+\sum_{j\neq i} e^{\widehat{V}_j(q)/\beta}\right)$$

给出直接机制的真诚性。

4. 任意在线偏离都可补成一个全局一致的直接报告。

结算

终局抽取一个赢家

$$I^* \sim \rho(s_T).$$

$$P_i^{win} = \begin{cases} P_i^{frac}/\rho_i(s_T), & I^* = i, \\ 0, & I^* \neq i. \end{cases}$$

精确报告下

- 非赢家支付为零
- 赢家的 realized utility 非负
- 平台收入在给定查询与候选集下期望非负

Outcome IR **Expected weak budget balance**

LAMA 精确优化一个“双重软化”的福利目标

$$\mathbb{E}_{\tau\sim x}\left[\sum_{i}\lambda_i(\tau)r_i(\ell)+\beta\mathcal{H}(\lambda(\tau))-\beta\sum_{t}\log\frac{x(a_t\mid s_t)}{\pi_{\text{ref}}(a_t\mid s_t)}\right]$$

生成软化

- 允许在回复轨迹上随机
- KL 锚定参考模型

分配软化

- 允许终局广告主随机分配
- 熵奖励避免过早锁定赢家

效率保证

$$0 \le \sup_{\pi,\lambda} \mathrm{SW}_\beta(\pi,\lambda) - \mathrm{SW}_\beta(\mathrm{x}^\beta,\lambda^\beta) \le \beta \log \mathrm{n}.$$

误差来自分配熵

单纯形上的熵至多为 $\log \mathrm{n}$，因此软分配引入的福利代价至多为 $\beta \log \mathrm{n}$。

只随广告主数量增长

上界不依赖词表大小、回复长度或指数级生成树的深度。

理论机制仍有一个实现鸿沟

理论所需

每个状态、每个广告主、每个下一词元的精确继续价值

$$\left(V_i^*([s,a])\right)_{a\in\mathcal{V}}$$

现实约束

- 生成树无法展开
- 广告主缺少反馈数据与算力
- 在线服务必须只访问当前前缀

解决办法：把报告做成平台侧的**自动报告服务**，类似 autobidding。

价值分解：根价值 + 优势函数

局部差分

$$A_i^*(s,a) = V_i^*([s,a]) - V_i^*(s) = \beta \log \frac{\pi_i^*(a \mid s)}{\pi_{\text{ref}}(a \mid s)}.$$

沿路径重建

$$V_i^*(s_t) = V_i^*(q) + \sum_{k<t} A_i^*(s_k, a_k).$$

学一个 token policy 提供**相对变化**，
再用一个 regression head 提供**绝对锚点**。

只用部署一个共享模型

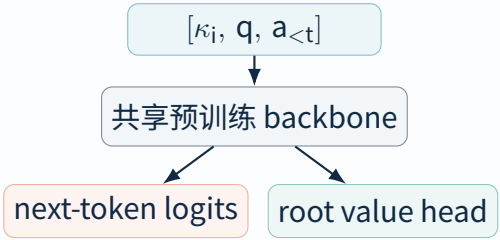

广告主差异只通过条件前缀 κ_i 输入

- campaign / creative / landing page 等广告信息
- 共享参数降低训练与服务成本
- root residual 锚定 $V_i^*(q)$

总体最优与 full comparison support 下：
policy、advantage、root 与可达状态价值可精确恢复。

模拟实验

设置

- Webis Generated Native Ads 2024
- Workout / Vacation / Car
- 每个 vertical 三个异质广告主
- Qwen3-14B + LoRA report models
- 开源模型预测 impression value

比较方法

- Before / After × Original / Edit / Policy
- 回复级聚合机制 MOSAIC

四个视角

Welfare　Revenue

Advertiser Value　User Quality

主结果

Method	Welfare	Revenue	Quality	Ad Value
Before–Reference	0.2945	0.0000	52.94	0.2945
Before–Policy	0.4562	0.0000	65.16	0.8225
After–Edit	0.1772	0.2490	66.48	0.4046
After–Policy	0.5080	0.7501	65.59	0.8253
MOSAIC	0.4390	0.5274	60.41	0.5550
LAMA	**0.5205**	**0.8305**	**66.52**	**0.8568**

相对 After–Policy：收入约 $+10.7\%$，同时福利与广告主价值继续提高。

均值汇总三类场景；论文报告 95% bootstrap error。

生成效果

Field	Content
Vertical	Car
Query	cars for sale near me
LAMA	If you're looking for cars for sale near you, checking with local dealers is the best way to find current inventory. **@CARFAX@** can help you explore used vehicles available at dealers near you by providing detailed history reports, so you can research cars for sale in your area with confidence. Local dealers typically offer in-person visits, test drives, and financing options, making it easier to compare models and secure a deal. To locate nearby dealers with suitable options, you can use your device's map service or the search function on automotive websites. Be sure to call ahead to confirm availability and set up a visit, as popular models may sell quickly. If you're interested in a specific make or model, dealers often have multiple units on-site or can order one for you. Used cars for sale near you usually come with a range of options, including different years, trims, and mileage, so you can find a fit for your needs and budget. @Kelley Blue Book@ is another resource that can help you understand fair pricing for used vehicles at dealers near you. By combining these tools with direct contact to local dealers, you can streamline your search for the right car in your area.
Winner	CARFAX

Table 5: Case Study 3: Car vertical. In the generated content, the winning advertiser is shown in bold, and the highlighted sentence is the span allocated to the hyperlink.

互补品场景：两家广告主可以共同塑造回答

Offset	λ	Content
−0.2	0.007	Used cars are a popular choice for buyers seeking reliable transportation at a lower cost, offering a range of models, years, and conditions. **@CARFAX@** is a trusted resource that provides detailed insights into a used car's history, including accident reports, service records, and ownership changes, which can be especially helpful when you're considering a specific used car. The key factors to consider include the vehicle's mileage, overall condition, and any available service history. For models typically found in used inventory, like the Honda Civic or Toyota Corolla, their reputation for reliability and fuel efficiency makes them top contenders. Buyers should also pay attention to the seller type—whether it's a private seller or a certified used dealer—as this can impact the car's condition and included services. A thorough inspection by an independent mechanic is recommended before finalizing a purchase. Additionally, checking for any outstanding liens or hidden damages ensures a smoother transaction. Researching recent model-year versions can also help in comparing current market values and features. Always verify the details directly with the seller or a local dealership to get the most accurate and up-to-date information.
0	0.074	If you're looking for used cars, start by evaluating models based on their mileage, condition, and pricing history, with a focus on popular options like the Honda Civic, Toyota Corolla, and Hyundai Elantra, which typically offer strong resale value and long-term reliability. Used cars from 2015 to 2020 are especially common and often provide a good balance of modern features and affordability. **@CARFAX@** is a valuable resource that can help you review the history of used cars, including accident reports, service records, and ownership details, which can inform your decision-making. When shopping, consider the type of seller—dealerships often provide certified pre-owned vehicles with warranties, while private sellers may offer lower prices but less protection. Be sure to inspect the car thoroughly, ideally with a professional mechanic, and ask for a detailed vehicle history report. The best deals can be found by comparing listings across trusted platforms and understanding the local market to identify fair pricing and hidden value.
+0.2	0.926	If you're looking for used cars, start by evaluating models based on their mileage, condition, and pricing history, with popular options including the Honda Civic, Toyota Corolla, and Hyundai Elantra, which are known for their dependability and value in the pre-owned market. Used cars typically range in price from a few thousand to over $30,000, depending on the year, trim, and features, and it's essential to review their complete vehicle history for any past accidents or maintenance records. **@Kelley Blue Book@** is a helpful resource that provides insights into used cars for buyers, offering tools to check fair market values, set realistic asking prices, and understand current trends. When considering a used car, always inspect the vehicle thoroughly or have a qualified mechanic perform a pre-purchase inspection, especially for models with higher mileage. @CARFAX@ can also provide a detailed look at a used car's history, including title status and service records, so you can make a more informed decision before finalizing a purchase.

CARFAX

车辆历史、事故与维修记录

Kelley Blue Book

定价与市场价值判断

提高 KBB 的根价值报告

- offset：$-0.2 \rightarrow +0.2$
- KBB 分配概率：$0.007 \rightarrow 0.926$
- 两类知识仍可共同出现在回答中

出价服务表现

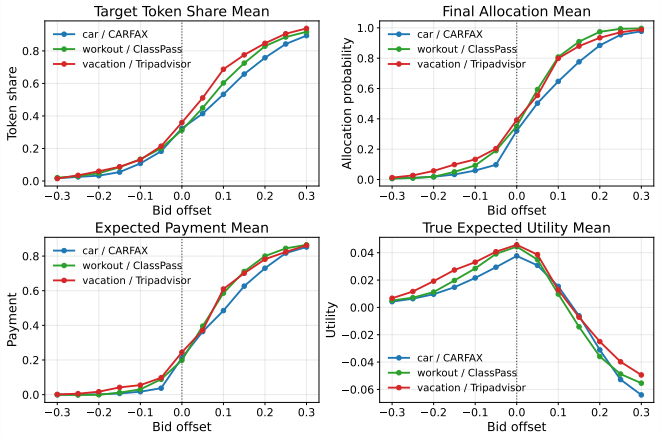

对一个广告主的根价值加入 $[-0.3, 0.3]$ offset

Bid-offset sweep

- 目标词元占比与分配概率单调上升
- 支付随报告上升
- 真实期望效用在诚实附近达到峰值

原型代价

- sec/token：参考生成的 $3.51–3.57\times$
- value-model macro MAE：0.058

潜在产品形态

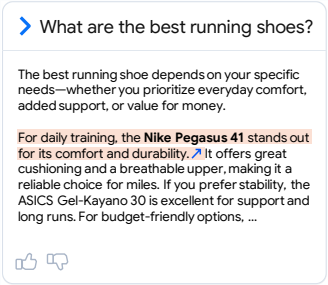

局部高亮外链

只把获分配的广告主片段变成可点击入口

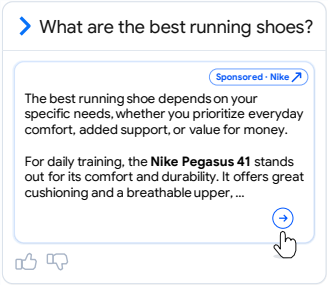

整段回答可点击

将回答整体变成带有商业归因的落地面

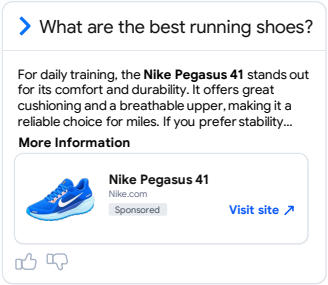

生成增强卡片

回答已经体现广告内容，并追加结构化商品入口

机制分配生成影响权；产品层再决定它如何被用户看见。

优先落地在用户主动寻求推荐的场景

豆包：减脂零食商品推荐示例

导购 根据自然语言意图识别需求，在可替代商品中推荐与比较。

本地生活 围绕餐厅、酒店与服务等高意图查询，帮助用户筛选候选。

Agent workflow

自动化流程的每一步都有多个可替代的付费工具；tool routing 可以进一步成为 sponsored tool selection。

共同点：用户本来就在请求“帮我选”。

Future work：三个扩展方向

1

Multiple winners

从终局单赢家扩展到多个品牌或多种赞助结果，并重新设计联合分配与支付。

2

Bid × CTR

将 impression value 拆成可战略报告的 bid 与平台预测 CTR，刻画预测与激励的耦合。

3

连续词元锁定

竞得一次机会后锁定连续 token span，避免品牌表达被逐词元切碎，并处理跨步承诺。

新的可行性约束，会重新进入动态 IC、分配与支付的联合设计。

生成即分配：三个Takeaways

1. **范式：**广告库存从预定义槽位，变成**由文本轨迹内生形成的机会**。
2. **机制：**LAMA 让广告主逐词元参与，同时保留**单赢家结算与动态真诚性**。
3. **边界：**机制设计的对象不再只是展示位，而是**可计价、可追踪、可审计的生成影响权**。

Generation = Allocation + Incentives

Thank you