Approximating Drift-Diffusion Models for User Decisions Under Nudging and External Information

Gustavo Grivol   Hanna Halaburda   Alex Tuzhilin

ICML 2026

Problem Visualization & Objective

- **History:** Displayed advertisements $X_i \geq 0$ (red) and user actions τ_i (blue).

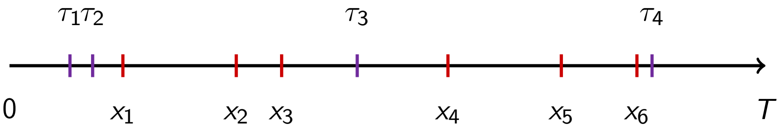

- **First Challenge: Modeling** — Estimate the probability of the next action given history $\mathcal{H}_t$ and signals X_t:

$$P\left(\tau_{next} \in (t, T] \middle| \mathcal{H}_t, X_t\right)$$

- **Second Challenge: Nudge Timing** — Optimize signal timing to maximize activity (e.g., Daily Active Users):

$$\sup_{0<x_1<\cdots<x_N<T} P\left(\tau_1 \in [0, T] \middle| \{x_i\}_{i=1}^{N}\right)$$

Foundation for Our Model

Requirements

- Reasonable spacing between actions (resting time).
- Signals are not immediately processed by users, inducing smoothness over the distribution).
- Be able to analytically separate the effect of the nudges.

Our Inspiration: In the neuroscience literature there is the Drift Diffusion Model (DDM), originally studied for the Two-Alternative Forced-Choice (2AFC). For single choice, we have the Go/No-Go:

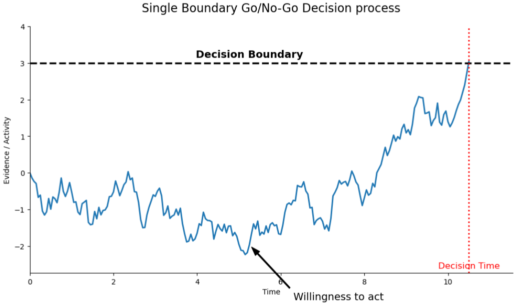

Dynamics of the Evidence Process

- Stochastic Differential Equation for Z_t (Willingness to act):

$$dZ_t = \mu(X(t), t)\, dt + dB_t$$

  Where dB_t is a standard Brownian motion, and $\mu(X(t), t)$ is the drift of the process, depending on the time and some general process of external information $X(t)$ (where we incorporate nudges, notifications, ...).
- Agents act when $Z_t = u$, then the evidence process resets: $Z_{t+} = 0$
- History notation: $\xi_t := (\tau_t, (X(s))_{s\in[0,t]})$ where τ_t is the last hitting time prior to t.

Question

Problem: If we introduce a general drift in the Brownian motion allowing incorporation of external information, the **model is not tractable**. Estimating μ over large classes of functions is extremely hard.

Main Theorem: Distribution Approximation

Theorem (Approximation for High Thresholds)

Consider the auxiliary function where the numerator represents the curved boundary:

$$f(t,u,\xi_t) := \frac{u - \int_{\tau_t}^{t} \mu(X(s),s)\,ds}{(t-\tau_t)^{1/2}}$$

If $\mu(X(t),t)$ is Riemann integrable and the drift is bounded $|\mu(X(t),t)| < M$, as $\mathbf{u} \to \infty$:

$$P(\tau_u^k \in [t,a] | \tau_t) \approx \int_t^a \left[\frac{f(s,u,\xi_s)}{2(s-\tau_t)} - f_t(s,u,\xi_s) \right] \phi(f(s,u,\xi_s))\,ds$$

where ϕ denotes the standard normal density and f_t is the partial derivative with respect to time.

This enables MLE and allows us to analyze the nudging optimization directly through the distribution. We will refer to the approximate model as **ExtDDM**.

Signal Optimization & Structure

Maximizaing

$$\max_{\tau_{last} \leq x_1 < \cdots < x_J \leq a} P(\tau_{next} \leq a \mid \tau_{last}, \{x_j\}_{j=1}^{J})$$

- **Scenario 1: Permanent Gain**: $\Psi(h(s)) \geq 0$ for all $s > 0$. Signals permanently push the process toward the boundary.
- **Scenario 2: Temporary Gain** $\Psi(h(s)) \geq 0$ for $s \in [0, r), \quad \Psi(h(s)) < 0$ for $s > r$

Proposition 1

Assume **Scenario 1**: For any $a > 0$, there exists $u^* > 0$ such that $x = 0$ maximizes the objective OPT_a.

Proposition 2

Assume **Scenario 2**: For all $a > 0$, there exists $u^* > 0$ such that if $u > u^*$, then the objective OPT_{a,τ^1} is strictly increasing in the signal position x, for all $x \in [0, a - r)$.

Main Empirical Application: KuaiRec 2.0 and Hybrid Modeling

Dataset Overview

- **Scale:** Approx. 12.3 million watch-time logs.
- **Features:** User demographics, video attributes, and historical viewing patterns.

The "Quick-Skip" Challenge

- **Behavior:** Users often switch immediately ($t \approx 0$) if content is not engaging.
- **Modeling Issue:** Many samples contain virtually no decision process.
- **Dual Mode:** Browsing (short views) vs. Consuming (long views).

Two-Stage Hybrid Model

1. **Gatekeeper (Stage 1):** Pre-trained EGMN predicts whether watch-time exceeds the 66% quantile.
2. **Predictor (Stage 2):**
   - *Short View:* Use EGMN prediction.
   - *Long View:* Use **ExtDDM**.

ExtDDM Architecture

- **Encoder (f_{enc}):** 2-layer NN + LayerNorm producing latent $Z_{i,j}$.
- **Threshold (f_{thr}):**

$$u(Z_{i,j}) = \ln\left(1 + e^{f_{\mathrm{thr}}(Z_{i,j}) + u_{\mathrm{const}}}\right).$$

Results

Baselines: VR (Wide&Deep), TPM (Tree-based), CREAD (Classification-Restoration), EGMN (Mixture).

Method	MAE	XAUC	KL Divergence
VR	4.6936	0.5383	1.4781
EGMN	4.4597	0.5805	0.8570
TPM	5.8098	0.5536	2.4928
CREAD	4.3863	0.5883	2.0255
ExtDDM + EGMN	**4.2833**	**0.6053**	**0.2063**
Improvement	*+2.3%*	*+16.1%*	*+75.9%*

Summary of Contributions

We provide a decision time approximation from a well grounded model and still leveraging:

1. **Closed-form approximation**
2. **Analysis over Optimal nudge timing**
3. **Strong Empirical validation**

Thank You!