

Normalized Conviction: The $2^{g/H}$ Combination Rule

Universal Model Project

March 2026

Abstract

We test three hybrid policies combining support-gap and entropy signals for n-gram order blending: (H1) $w = \exp(-H) \cdot 2^g$, (H2) gap-weighting restricted to low-entropy orders, and (H3) $w = 2^{g/H}$ where g is the support gap and H is the entropy. H3 is the strongest higher-order hybrid across the tested orders: it remains below ent-blend at order 4, but improves on pure gap there, then beats gap-blend at orders 6 and 8 with an advantage that grows from 0.005 bpc (order 6) to 0.105 bpc (order 8). The gap/entropy ratio measures *normalized conviction*: how committed a predictor is, relative to how many options its distribution supports.

1 Motivation

The conviction–accuracy tradeoff (see companion note) showed that gap-blend wins at orders ≥ 6 despite worse selection accuracy, because its extreme weighting produces large per-position wins. But at order 4, ent-blend still wins by 0.085 bpc: the gap signal is too noisy at low orders (many contexts unseen, gaps dominated by unigram). Can we combine both signals?

Three candidate hybrid policies:

H1. $w_k = \exp(-H_k) \cdot 2^{g_k}$ (product of both)

H2. $w_k = 2^{g_k}$ if $H_k < \text{median}(H)$, else 0 (entropy-filtered gap)

H3. $w_k = 2^{\lfloor g_k/H_k \rfloor}$ (gap normalized by entropy)

2 Results

Order	gap-blend	ent-blend	H1	H2	H3
4	2.509	2.424	2.463	4.630	2.479
6	2.244	2.337	2.367	5.632	2.239
8	2.294	2.379	2.381	6.528	2.189

H3 advantage over the best previous policy at each order:

Order	Best prev	H3	Improvement
4	2.424 (ent)	2.479	−0.055 (worse)
6	2.244 (gap)	2.239	+0.005
8	2.294 (gap)	2.189	+0.105

At order 4, H3 is worse than ent-blend but better than gap-blend. At orders 6 and 8, H3 is the outright winner.

3 Why H3 Works

3.1 The gap/entropy ratio

The support gap $g = s_1 - s_2$ measures absolute commitment: the difference between the top two log-probabilities. The entropy H measures how many effective options the distribution supports.

The ratio g/H measures *commitment per unit of distributional complexity*:

- High g/H : the distribution is both committed *and* concentrated. Strong signal.
- High g , high H : committed on the top event but spread out below. Weaker signal than the gap suggests.
- Low g , low H : concentrated but not committed to a specific winner. The distribution is sharp but contested.

By normalizing gap by entropy, H3 avoids the failure mode of pure gap-blend: over-weighting low orders (especially unigram) whose high gap comes from a narrow mode in an otherwise flat distribution.

3.2 Why H1 and H2 fail

H1 (product) gives extreme weight to orders that are *both* concentrated and gap-dominant, which is too selective: it discounts orders with moderate signals that still carry useful information.

H2 (entropy filter) discards all orders above the median entropy, losing mass that the blend needs for fall-through coverage. It is catastrophic because early in training, most orders have high entropy, and filtering them out leaves only the unigram.

3.3 Scaling behavior

The advantage of H3 over pure gap grows with order because higher orders have more spurious gap signals from newly created n-gram contexts. A context seen twice (counts [1,1]) has $g = 0$ and $H \approx 1$, giving $g/H \approx 0$ —correctly assigned low weight. The same context under pure gap-blend might still compete because $2^0 = 1$ is non-negligible when few orders are active.

4 Conclusion

A strong current candidate combination rule for higher-order UM n-gram models is $w_k = 2^{\lfloor g_k/H_k \rfloor}$, the *normalized conviction* policy. It is not the winner at order 4, but it is the strongest hybrid across orders 4–8 and its advantage grows with order. It remains a native tropical operation (2^x is a shift) and can be implemented with no additional parameters.

5 Reproduce

```
./umr hybrid-blend enwik9 4 1000000
./umr hybrid-blend enwik9 6 1000000
./umr hybrid-blend enwik9 8 1000000
```