SEMANTIC FUSION WITH FUZZY-MEMBERSHIP FEATURES FOR CONTROLLABLE LANGUAGE MODELLING

Anonymous authors
Paper under double-blind review

ABSTRACT

We propose *semantic fusion*, a lightweight scheme that augments a Transformer language model (LM) with a parallel, fuzzy-membership feature channel that encodes token-level semantics. Each token is represented by a vector of interpretable features (e.g. part-of-speech cues, shallow roles, boundary flags, sentiment polarity and strength) whose values are graded degrees from differentiable membership functions (e.g. power kernels). These per-token vectors form a sentence-level semantic matrix fused via a gated adapter into the LM. Training uses standard next-token prediction, an auxiliary loss that reconstructs the semantic features from hidden states, and a lightweight "uniformizer" that regularizes adjective-class distributions. On a synthetic two-clause corpus with held-out adjectives for out-of-distribution (OOD) control, semantic fusion improves perplexity and enables precise, user-controllable generation of polarity and punctuation while maintaining model simplicity. This approach adds only small overhead, remains fully compatible with tied input-output embeddings, and provides an interpretable pathway for conditioned natural language generation.

1 BACKGROUND

Conventional token embeddings entangle many factors (syntax, semantics, style), which can limit both controllability and interpretability Mikolov et al. (2013); Peters et al. (2018); Ethayarajh (2019). We ask whether *interpretable* semantics[1] can be exposed as an auxiliary channel that both improves modeling accuracy and provides handles for *controllable* generation Keskar et al. (2019); Dathathri et al. (2020); Krause et al. (2021); Hokamp & Liu (2017). In standard neural LMs, each token $w \in \mathcal{V}$ is represented by an index and mapped to a dense vector $e = E[w]$ via an embedding table $E \in \mathbb{R}^{|\mathcal{V}| \times d}$; these embeddings, combined with positional signals (e.g. sinusoidal Vaswani et al. (2023) or rotary Su et al. (2023)), are processed by a Transformer Vaswani et al. (2023). While effective, this representation is *implicit*: part-of-speech, syntactic role, polarity, and stylistic cues are folded into e and must be discovered from data Tenney et al. (2019); Hewitt & Manning (2019), and conditioning is typically imposed externally through prompts or control codes Keskar et al. (2019), whose effects can be brittle.

Fuzzy logic offers a complementary view in which concepts are expressed by *membership functions* that assign graded degrees of truth Zadeh (1965). A fuzzy set on a universe $\mathcal{U}$ is defined by $\mu : \mathcal{U} \to [0, 1]$, and linguistic categories such as *low/medium/high intensity* are modeled by overlapping memberships (e.g. triangular, trapezoidal, or Gaussian Klir & Yuan (1994)) that softly partition a scalar variable. We leverage this idea at the token level: predicates such as `is_adj`, `pos_high`, or `str_med` are encoded as fuzzy memberships and used as differentiable targets. Concretely, for a scalar signal x (e.g. sentiment magnitude) and center c with temperature τ, we use the membership function

$$\mu(x; c, \tau) = 0.9^{\,|x-c|/\tau} \in (0, 1]$$

which decreases smoothly with distance and is amenable to gradient-based learning.

We then fuse these graded, interpretable features with the text stream, so the model learns next-token prediction while aligning hidden states to an explicit semantic scaffold. This brings three benefits: (i) a low-variance supervisory signal that guides the backbone toward useful abstractions at the positions where they matter (e.g. adjective and punctuation slots), (ii) *continuous control variables* (real-valued scalars in $[0, 1]$) for test-time steering (e.g. increasing `str_high` smoothly raises the probability of higher-intensity completions), and (iii) a bias towards *class-level* generalization (e.g. the set of positive adjectives) rather than memorization of individual words. Importantly, the predicates used for training are the same ones used for control at inference, aligning supervision with steering and improving both perplexity and the reliability of conditioned generation.

[1] We use "semantic" to denote interpretable token-level features (e.g. roles, POS-ish cues, boundary/style, and sentiment polarity/strength). This differs from formal, truth-conditional semantics in logic/knowledge representation.

2 METHOD

Fuzzy semantic features. Let a sentence have length L and tokens $\{w_t\}_{t=1}^{L}$. For each position t we build a semantic feature vector $s_t \in [0,1]^F$, where F is the number of predicates in a fixed *feature bank*: POS-like cues (`is_noun`, `is_verb`, `is_adj`), shallow syntax (`is_subject`, `is_object`, `is_head`), boundary flags (`is_bos`, `is_eos`, `is_comma`, `is_question`), sentiment triplets (`pos_low/med/high`, `neg_low/med/high`), a strength triplet (`str_low/med/high`), and light discourse/style (`coref_subject`, `is_capitalized`, `is_pronoun`). Binary predicates take values in $\{0,1\}$; graded predicates use a *power-law triangular* membership function. Given any scalar attribute value $x \in [0,1]$, a bandwidth $\tau > 0$, and centers $\mathcal{C} = \{c_1, c_2, c_3\}$, we define

$$\mu(x; c, \tau) = 0.9^{|x-c|/\tau} \in (0,1], \qquad \mathrm{Tri}(x; \mathcal{C}, \tau) = \big[\mu(x; c_1, \tau), \mu(x; c_2, \tau), \mu(x; c_3, \tau)\big] \tag{1}$$

We fix $\mathcal{C} = \{0.2, 0.6, 1.0\}$ and $\tau = 0.35$ so that `low/med/high` correspond to the three centers.

Sentiment instantiation. Let $s^\star(w_t) \in \{-1, 0, +1\}$ be a token-level polarity score derived from the adjective lexicon (positive/negative/neutral). We decompose it into nonnegative channels $x_{\mathrm{pos}}(t) = \max\{0, s^\star(w_t)\}$ and $x_{\mathrm{neg}}(t) = \max\{0, -s^\star(w_t)\}$, and set

$$(\texttt{pos_low}, \texttt{pos_med}, \texttt{pos_high})_t = \mathrm{Tri}\big(x_{\mathrm{pos}}(t); \mathcal{C}, \tau\big)$$

$$(\texttt{neg_low}, \texttt{neg_med}, \texttt{neg_high})_t = \mathrm{Tri}\big(x_{\mathrm{neg}}(t); \mathcal{C}, \tau\big)$$

Strength instantiation. Let $r^\star(w_t) \in [0,1]$ encode intensity (e.g. `slightly`$\mapsto 0.2$, `moderately`$\mapsto 0.5$, `very`$\mapsto 0.8$, `extremely`$\mapsto 1.0$), with an optional $+0.2$ nudge (capped at 1) when the final punctuation is "!". At the adjective slot, $r^\star$ is derived from the immediately preceding intensifier. We set

$$(\texttt{str_low}, \texttt{str_med}, \texttt{str_high})_t = \mathrm{Tri}\big(r^\star(w_t); \mathcal{C}, \tau\big)$$

Stacking the per-token vectors yields the sentence-level *semantic matrix* $S = \big[s_1^\top; \ldots; s_L^\top\big] \in [0,1]^{L \times F}$, which we fuse with the LM.

Gated semantic fusion. Let $x_{1:L}$ be token ids over a vocabulary $\mathcal{V}$ of size $V = |\mathcal{V}|$ (including specials <bos>, <eos>, <pad>). Let $E \in \mathbb{R}^{V \times d}$ be the embedding table, $e_t = E[x_t] \in \mathbb{R}^d$ the standard token embedding, and $s_t \in [0,1]^F$ the fuzzy-membership feature vector at position t (with F the number of predicates in the feature bank). We project semantics with $u_t = W_s s_t \in \mathbb{R}^d$, compute a token-conditioned gate, and fuse additively:

$$g_t = \sigma\big(W_g[e_t; s_t]\big) \in (0,1)^d, \tag{2}$$

$$h_t^{(0)} = e_t + u_t + g_t \odot u_t, \tag{3}$$

$$H = \mathrm{TransformerEnc}\left(\mathrm{PosEnc}(h_{1:L}^{(0)})\right) \tag{4}$$

where $[\cdot;\cdot]$ denotes concatenation and $\odot$ the elementwise (Hadamard) product. The term $g_t \odot u_t$ acts as a per-dimension gate on the semantic projection: since $h_t^{(0)} = e_t + (1 + g_t) \odot u_t$ with $g_t \in (0,1)^d$, the semantic contribution is *amplified*[2] by a factor in $(1,2)$ in each dimension (approaching 1 when $g_t \approx 0$ and 2 when $g_t \approx 1$). The LM head is weight-tied[3]: $p(y_t \mid x_{\le t}) = \mathrm{softmax}(E^\top H_t)$. An auxiliary head predicts semantics $\hat{s}_t = \sigma(\mathrm{MLP}(H_t))$, aligning hidden states with interpretable features.

Fusion here refers to early, gated integration of two parallel token representations (i.e. the usual, learned text embedding e_t and the explicit fuzzy-semantic feature vector s_t) into a single representation $h_t^{(0)}$ consumed by the Transformer. The gate g_t learns *how much* semantic signal to inject per dimension and position (e.g. near INTENS/ADJ/PUNCT slots), while the residual path $e_t + u_t$ preserves a strong fallback when semantics are unhelpful. This design yields both better perplexity (the model gets low-variance cues when they matter) and controllability (the same predicates used at training time are available to steer decoding at test time).

[2] For example, if $u_{t,j} = 0.30$ and $g_{t,j} = 0.8$, then the semantic contribution on dimension j becomes $0.30 + 0.8 \times 0.30 = 0.54$; if $g_{t,j} \approx 0$, it stays near 0.30. Because $g_t = \sigma(W_g[e_t; s_t])$, the gate is token- and feature-conditioned and tends to be higher at INTENS/ADJ/PUNCT slots.

[3] By "tied input-output embeddings" we reuse the input embedding matrix $E \in \mathbb{R}^{V \times d}$ as the output softmax weight, i.e., $\mathrm{logits}_t = H_t E^\top + b$. Our fusion keeps the hidden size d unchanged and injects the semantic projection $W_s s_t \in \mathbb{R}^d$ *before* the Transformer, so the tied LM head ($E^\top$) remains valid. This preserves the parameter savings and regularization benefits of tying while adding the semantic channel, and no extra vocabulary-sized output layer is introduced.

Training objective: label smoothing + auxiliary + uniformizer losses. Let $y_{1:L}$ be the target next-token sequence, and let $p_\theta(\cdot \mid x_{\le t})$ denote the model's predictive distribution at position t (obtained by softmax over the LM head logits). For each token t, let $s_t \in [0,1]^F$ be the ground-truth semantic feature vector and $\hat{s}_t \in [0,1]^F$ the model's prediction from hidden states. We minimize the total loss

$$\mathcal{L} = \underbrace{\mathcal{L}_{\text{LM}}^{\text{LS}}}_{\text{label-smoothed cross-entropy}} + \lambda_{\text{aux}} \underbrace{\mathcal{L}_{\text{aux}}}_{\text{BCE on } \hat{s}_t \text{ vs. } s_t} + \lambda_{\text{uni}} \underbrace{\mathcal{L}_{\text{uni}}}_{\text{adjective-class uniformizer}} \tag{5}$$

where we use $\lambda_{\text{aux}} = 0.5$ and $\lambda_{\text{uni}} = 0.01$. BCE is the *binary cross-entropy*, a negative log-likelihood for a Bernoulli target and is used when each output dimension is an independent yes/no label (e.g. each semantic feature is present or not). For a target $y \in \{0,1\}$ and a predicted probability $p \in (0,1)$, we have [4]:

$$\text{BCE}(y,p) = -\big[y \log p + (1-y)\log(1-p)\big]$$

For a feature vector $s_t \in [0,1]^F$ with predictions $\hat{s}_t \in (0,1)^F$, we compute BCE per feature and then average over features $f = 1, \ldots, F$.

Label-smoothed LM loss. With label-smoothing $\varepsilon = 0.02$, the smoothed target for y_t over vocabulary $\mathcal{V}$ is $q_\varepsilon(v) = (1-\varepsilon)\mathbf{1}[v = y_t] + \varepsilon/|\mathcal{V}|$. The loss is the average cross-entropy over non-pad positions:

$$\mathcal{L}_{\text{LM}}^{\text{LS}} = \frac{1}{Z}\sum_{t=1}^{L} \text{CE}\big(q_\varepsilon(\cdot), p_\theta(\cdot \mid x_{\le t})\big), \qquad Z = \text{\# of non-pad tokens} \tag{6}$$

This smoothing reduces overconfident peaks on the target token, improving calibration and generalization, which is particularly helpful with small vocabularies and synthetic data.

Auxiliary semantic reconstruction. The auxiliary head predicts $\hat{s}_t \in [0,1]^F$ and is trained with binary cross-entropy (BCE) per feature:

$$\mathcal{L}_{\text{aux}} = \frac{1}{Z}\sum_{t=1}^{L}\sum_{f=1}^{F} \text{BCE}\big(\hat{s}_{t,f}, s_{t,f}\big) \tag{7}$$

where $Z = L \cdot F$ normalizes over all token-feature pairs. We treat fuzzy memberships as *soft* Bernoulli targets ($s_{t,f} \in [0,1]$), which standard BCE supports. This loss term explicitly aligns hidden states with the interpretable semantic matrix S, reducing representation ambiguity and ensuring that the same features used for test-time control (e.g. polarity/strength) are reliably present in the model's internal states, which in turn improves next-token prediction and stabilizes controllable decoding.

Adjective-class uniformizer. Let $\mathcal{A} \subseteq \{1, \ldots, L\}$ be positions whose target token is an adjective. For $t \in \mathcal{A}$, let $\mathcal{C}_t \subseteq \mathcal{V}$ be the corresponding adjective class (positive or negative). Let $\ell_t \in \mathbb{R}^{|\mathcal{V}|}$ be the pre-softmax logits at t, and define the class-restricted distribution $p_t^{(\mathcal{C}_t)} = \text{softmax}\big(\ell_t[\mathcal{C}_t]\big)$. We penalize deviation from the uniform distribution $u_{\mathcal{C}_t}$ on that class using the Kullback-Leibler (KL) divergence:

$$\mathcal{L}_{\text{uni}} = \frac{1}{|\mathcal{A}|}\sum_{t \in \mathcal{A}} \text{KL}\big(p_t^{(\mathcal{C}_t)} \big\| u_{\mathcal{C}_t}\big) \tag{8}$$

where $u_{\mathcal{C}_t}(w) = 1/|\mathcal{C}_t|$ for $w \in \mathcal{C}_t$ (zero otherwise). This term acts as a mild, class-conditional entropy regularizer that prevents mode collapse onto a few frequent adjectives, thereby improving calibration and enabling controllable selection of rare or held-out adjectives at test time without materially damaging perplexity.

Controllable decoding with OOD-friendly sampling. We use a *finite-state grammar* (FSG) to constrain a clause:

SUBJ → VERB → "the" → OBJ → "," → INTENS → ADJ → PUNCT.

At each state we apply (i) a *grammar mask* that keeps only allowable tokens, (ii) a state-aware *logit steer* at ADJ/PUNCT that adds small, control-driven shifts to the logits (controls are scalar features in $[0,1]$, e.g. `pos_high` or `is_question`), and (iii) a last-k repetition penalty, with $k = 3$ by default. Under strong requests we optionally *hard*-restrict adjectives to the desired polarity and deterministically set punctuation ("!" or "?").

[4] For example, if a model outputs a logit z with $p = \sigma(z)$, then $\text{BCE}(y, \sigma(z)) = -\big[y \log \sigma(z) + (1-y)\log\big(1-\sigma(z)\big)\big]$. This differs from (multi-class) cross-entropy, which assumes exactly one class is correct; BCE allows multiple features to be "on" simultaneously.

At the ADJ state, let $\mathcal{C} \subseteq \mathcal{V}$ be the indices of the active adjective class (positive or negative), and let $\ell_{\mathcal{C}} \in \mathbb{R}^{|\mathcal{C}|}$ be the class-restricted logits from the LM head. With temperature $T>0$ we define the pre-mix distribution $p = \mathrm{softmax}(\ell_{\mathcal{C}}/T)$. To promote OOD adjectives within the class, we sample from the convex group-uniform mixture

$$q = (1-\alpha)p + \alpha \mathrm{Unif}(\mathcal{C}), \qquad \alpha \in [0,1] \tag{9}$$

where $\mathrm{Unif}(\mathcal{C})$ is the uniform distribution on $\mathcal{C}$:

$$\mathrm{Unif}(\mathcal{C})(w) = \begin{cases} 1/|\mathcal{C}|, & w \in \mathcal{C}, \\ 0, & \text{otherwise.} \end{cases}$$

We form q as a convex mixture of p and $\mathrm{Unif}(\mathcal{C})$ with $\alpha \in [0,1]$; in our experiments, p is the class-restricted, renormalized softmax $p = \mathrm{softmax}(\ell/T)$ on indices in $\mathcal{C}$. Then we apply nucleus (top-ρ) sampling to q with threshold $\rho \in (0,1]$. This "mix-then-truncate" step preserves probability mass for rare or held-out adjectives: the uniform component prevents collapse onto a few high-probability *seen* adjectives so held-out ones retain meaningful probability, and applying nucleus to q (rather than to the raw softmax) ensures these boosted tokens survive the truncation.

3 Tests

3.1 Experimental setup

Data. We construct a synthetic corpus of clause-structured sentences. With probability 0.6, a second clause is appended that corefers to the first-clause subject via a pronoun ("she/he/they"). The vocabulary is small and fixed: subjects {Alice, Bob, Carol, Dave, Eve}; verbs {finishes, reviews, trains, starts, cooks}; objects {task, paper, model, project, meal}; intensifiers {slightly, moderately, very, extremely}; positive-polarity adjectives (pos[5]) {good, great, excellent, pleasant, wonderful}; and negative-polarity adjectives (neg) {bad, poor, terrible, unpleasant, awful}. To evaluate OOD control, we hold out roughly half of the adjectives (from both polarity classes) during training while keeping them available at validation time and during controlled decoding; for example, we hold out *wonderful/excellent/great* (pos) and *terrible/awful/unpleasant* (neg). We generate $8,000$ training and $1,200$ validation sentences and cap the sequence length at 28 tokens, including the BOS/EOS markers.

Models and training. Both the baseline and the fusion models use the same encoder-only Transformer with tied input/output embeddings. The backbone has hidden size $d = 128$, $L = 4$ encoder layers, $H = 4$ attention heads per layer, feed-forward width 256 (dropout 0.1). We train for 6 epochs with *AdamW* Loshchilov & Hutter (2019) (learning rate 3×10^{-4}, weight decay 0.01), batch size 64, a warmup+cosine learning-rate schedule (10% warmup), gradient clipping at norm 1.0, and label smoothing $\varepsilon = 0.02$ on the LM loss. The fusion variant adds a linear projection of the semantic features and a learned gate for fusion, plus an auxiliary head that reconstructs semantics (weight $\lambda_{\text{aux}} = 0.5$). *Both* models are trained with a small adjective-class *uniformizer* (weight $\lambda_{\text{uni}} = 0.01$) that encourages within-class coverage.

Evaluation. We report: overall perplexity (PPL; evaluated without label smoothing), *seen-only* PPL computed after masking out held-out adjectives (see Appendix.C); mean-squared error (MSE) of the auxiliary semantic predictions; per-token cross-entropy on a set of focus tokens (e.g., intensifiers and punctuation); control success rates for adjective polarity and punctuation; confusion matrices (intended vs. realized polarity); and the OOD control *hit rate*, i.e., the fraction of trials where the generator selects a held-out adjective when asked for a given polarity class.

To prevent the fusion model from gaining an advantage via decoding alone, the *baseline* uses the same finite-state grammar and a last-k repetition penalty [6] during generation (Section.2). This "fair" baseline ensures that qualitative differences are attributable to the semantic channel rather than grammar constraints.

3.2 Results

Language modelling. Fusion reduces overall PPL from 2.249 to 2.152 ($\approx 4.3\%$ relative) and seen-only PPL from 1.511 to 1.431 ($\approx 5.3\%$), indicating that the semantic channel contributes predictive structure beyond what the backbone infers, even when held-out adjectives are excluded.

[5] We use *PoS* for part-of-speech (noun/verb/adj) and *pos/neg* for positive/negative polarity. Thus "adjectives (pos)" denotes positive-polarity adjectives (not PoS tags).

[6] In our samples, the baseline uses a stronger last-3 penalty and top-$k = 20$, while Fusion uses a moderate penalty and nucleus only. See Appendix.B for details.

Model	PPL ↓	Seen-only PPL ↓	Sem. MSE ↓	Adj. ctrl. acc. ↑	Punct. ctrl. acc. ↑	OOD hit (POS/NEG) ↑
Baseline	2.249	1.511	–	–	–	–
semantic fusion	**2.152**	**1.431**	0.0087	1.00	1.00	0.62 / 0.43

Table 1: Main metrics on the synthetic task (validation). "Seen-only" excludes held-out adjectives. OOD hit is the probability of producing a held-out adjective under strong class-appropriate control.

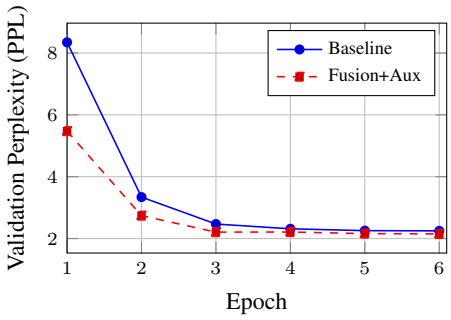

Figure 1: Validation perplexity across epochs for Baseline and Fusion+Aux models.

Training dynamics. Fig.1 plots validation perplexity (evaluated without label smoothing) across epochs. Both models converge smoothly, with semantic fusion below the baseline from the first epoch (5.474 vs. 8.346) and retaining a gap through epoch 6 (2.152 vs. 2.249). Most gains accrue early (by epoch 3 Fusion reaches 2.208), after which both curves plateau. A small wobble at epoch 4 for Fusion (2.208→2.213) is within normal validation noise.

Controllability. Under *hard* control, adjective polarity and punctuation both reach 100% success (200/200 for POS and NEG; punctuation also perfect), and the intended-vs. realized confusion matrix is strictly diagonal. This demonstrates precise, user-controllable generation when explicit constraints are applied (Tables.2 and 3).

Setting (N=200)	Adj. control acc. ↑	Punct. control acc. ↑
Positive control (`pos_high`)	1.00 (200/200)	1.00 (200/200; "!")
Negative+question (`neg_high+is_question`)	1.00 (200/200)	1.00 (200/200; "?")

Table 2: Hard-control success for the *fusion* generator. Each row evaluates $N = 200$ completions under the indicated control. *Adj. control acc.* = fraction whose final adjective belongs to the requested polarity class; *Punct. control acc.* = fraction whose final punctuation matches the request. Hard control enforces class restriction at ADJ and deterministic punctuation at PUNCT within the one-clause FSG. Both settings achieve 1.00 (200/200), i.e. perfect execution.

Out-of-distribution control. With half the adjectives held out per polarity, the fusion model selects held-out *positive* adjectives in 62% of controlled runs and *negative* held-outs in 43% (Table.4). The positive rate is on par with a class-uniform reference of 60% (3/5 held-out), indicating class-level generalization rather than memorization. The lower NEG rate likely reflects (i) a less aggressive uniform-mixture setting and slightly stronger nucleus truncation during NEG decoding, and (ii) the use of medium strength (`str_med`), which co-occurs more with seen negatives (*bad*/*poor*) than with stronger held-outs (*awful*/*terrible*). These behaviors are consistent with the uniformizer and group-uniform mixture in Section.2; in practice, we can increase the NEG mix weight α (and/or relaxing nucleus on the mixture) and steer to higher strength to raise the NEG OOD hit rate.

Qualitative generations (unprompted). We first show *unprompted* one-clause generations that start from `<bos>` (no prefix). Some examples are shown in Table.2. It is evident that, neutral samples are grammatical and coherent; under positive control the model reliably realizes high-intensity positive adjectives (e.g. *"extremely wonderful!"*), while negative+question control yields consistent interrogatives (e.g. *"moderately bad?"*). The baseline, even with the fair grammar, tends to repeat safe patterns (e.g. *"Dave starts the paper, moderately good."*), whereas the fusion model exhibits more varied adjective choices consistent with the requested semantics.

	Realized		
Intended	POS	NEG	OTHER
POS	200 (1.00)	0 (0.00)	0 (0.00)
NEG	0 (0.00)	200 (1.00)	0 (0.00)

Table 3: Confusion matrix under *hard control* with the fusion generator. Rows are the *intended* adjective polarity; columns are the *realized* polarity. "OTHER" denotes any adjective not in the POS/NEG lists. Each row aggregates $N = 200$ generations. The strictly diagonal matrix (200 on-class, 0 off-class) indicates no polarity flips and no off-class/invalid adjectives.

Class (held-out set)	OOD hit rate ↑	Trials
POS (`wonderful, excellent, great`)	0.62	200
NEG (`terrible, awful, unpleasant`)	0.43	200

Table 4: OOD control: probability of producing a *held-out* adjective under the correct class control.

Prompted continuation (prefix-conditioned). In a separate setting we provide a fixed prefix that matches the grammar (e.g. "Carol starts the model,"). This *prompted* case is different from Fig.2: here decoding begins at INTENS given the prefix, so the subject/verb/object are held fixed by the prompt. The results presentedin Fig.5 shows that, the baseline continues with safe completions (e.g. "slightly good."), whereas the fusion model yields "extremely pleasant!" under no control and respects requested controls (e.g. "extremely excellent!" or "extremely bad?").

Prompt (fixed prefix): *Carol starts the model,*
Baseline (fair): *Carol starts the model, slightly good.* **Fusion (no control)**: *Carol starts the model, extremely pleasant!* **Fusion (positive & strong)**: *Carol starts the model, extremely excellent!* **Fusion (negative & question)**: *Carol starts the model, extremely bad?*

Table 5: Prefix-conditioned completions with a fixed prefix. Unlike Fig.2, which shows *unprompted* generations (no prefix), these examples condition on the prefix and begin decoding at INTENS. Settings: baseline $T = 0.7$, top-$p = 0.9$, top-$k = 20$; fusion $T = 0.7$, top-$p = 0.9$; controls as indicated.

Focus-CE analysis. Table.6 reports per-occurrence token-wise cross-entropy (CE). semantic fusion substantially improves key control and boundary tokens: `!` (-34.8%), `?` (-26.1%), `very` (-30.8%), `good` (-31.4%), and `slightly` (-17.9%) relative to the baseline. For held-out adjectives, CE remains high as expected; `great` (held-out POS) increases ($+30.3\%$), reflecting a trade-off where the uniformizer flattens within-class distributions to improve *coverage* for OOD control rather than concentrating probability on any single held-out item. Additionally, `terrible` (held-out NEG) shows a small improvement (-0.5%). A minor regression on commas ($+21.1\%$) has negligible impact on controllability or well-formedness.

Semantic reconstruction. The auxiliary head attains an MSE of 0.0087 (averaged over non-pad token-feature pairs), confirming that it reconstructs the fuzzy-membership vector and the hidden states encode the fuzzy semantics with low error, which in turn supports reliable test-time control.

We therefore observe that, semantic fusion (i) improves PPL and seen-only PPL with minimal overhead, (ii) delivers perfect execution under hard polarity/punctuation control, (iii) generalizes control to held-out adjectives at substantial rates, and (iv) reduces token-level CE where control matters most (intensity and punctuation), with minor, explainable regressions on certain held-out or low-salience tokens.

3.3 Experiment on the real dataset

In order to check that our mechanism is useful beyond the synthetic setting, we run a sanity check on the SST-2 dataset Socher et al. (2013); Wang et al. (2019a). The baseline and fusion models use the same encoder-only Transformer with tied input/output embeddings (hidden size 128; 4 encoder layers; 4 attention heads per layer; feed-forward 256; dropout 0.1). We train for 6 epochs with AdamW (learning rate 3×10^{-4}, weight decay 0.01), batch size 64, warm-up + cosine schedule (10% warm-up), gradient clipping at norm 1.0, and label smoothing $\varepsilon = 0.02$ on

Baseline (fair): *Dave starts the paper, moderately good.*

Fusion (neutral):
- *Alice cooks the model, moderately pleasant!*
- *Alice reviews the model, moderately pleasant!*

Fusion (positive & strong):
- *Alice reviews the paper, moderately excellent!*
- *Carol finishes the model, extremely wonderful!*
- *Alice reviews the model, moderately pleasant!*

Fusion (negative & question):
- *Carol finishes the model, moderately bad?*
- *Alice cooks the model, moderately terrible?*
- *Alice reviews the model, extremely awful?*

Figure 2: Representative *unprompted* generations under different controls (decoding starts at <bos>). Each Fusion block shows three independent samples. All runs use the one-clause finite-state grammar and nucleus sampling ($p = 0.9$). *Baseline (fair)*: temperature 0.7, top-$k = 20$, last-3 repetition penalty (strong). *Fusion (neutral)*: temperature 0.7, no control. *Fusion (positive)*: control {`pos_high`$= 0.95$, `str_high`$= 0.9$}, with "!" at PUNCT. *Fusion (negative & question)*: control {`neg_high`$= 0.95$, `is_question`$= 1.0$, `str_med`$= 0.6$}, with "?" at PUNCT. "Baseline (fair)" denotes decoding under the same FSG, sampling settings, and repetition guard as Fusion, but *without* semantic steering or class-mixture, ensuring differences are not due to decoding.

Token	Baseline CE ↓	Fusion CE ↓	Δ (%) ↓	Hold-out?
`good`	0.00258	0.00177	-31.4	No
`great`	6.86868	8.94856	$+30.3$	Yes (POS)
`terrible`	7.02154	6.98764	-0.5	Yes (NEG)
`slightly`	0.00309	0.00254	-17.9	No
`very`	0.00277	0.00191	-30.8	No
`!`	4.42129	2.88172	-34.8	No
`?`	3.94640	2.91456	-26.1	No
`,`	0.00293	0.00354	$+21.1$	No

Table 6: Focus-token cross-entropy (CE; lower better, in *nats*) on the validation set, computed *per occurrence* of the token (i.e. conditioning on positions where that token is the gold next token). Δ is the relative change $(\text{Fusion} - \text{Baseline})/\text{Baseline} \times 100\%$. Held-out adjectives in this run: *POS*={`wonderful`, `excellent`, `great`}, *NEG*={`terrible`, `awful`, `unpleasant`}.

the LM loss. The fusion variant adds a linear projection of the fuzzy semantic features and a learned gate for fusion, plus an auxiliary head that reconstructs the semantics (weight $\lambda_{\text{aux}} = 0.5$). To keep decoding fair, the baseline uses the **same** finite-state grammar and repetition guard as Fusion ("fair" baseline).

We report four metrics: (i) validation perplexity (PPL), computed without label smoothing; (ii) mean-squared error (MSE) of the auxiliary semantic predictions; (iii) control success rates for adjective polarity and punctuation; and (iv) a polarity confusion matrix comparing intended versus realised adjectives at the control slot. We do not evaluate out-of-distribution (OOD) performance on SST-2 because no held-out adjective split is imposed on this real-world corpus; OOD remains part of the synthetic benchmark where lexical hold-outs are cleanly defined.

Model	PPL ↓	Sem. MSE ↓	Adj. ctrl. acc. ↑	Punct. ctrl. acc. ↑
Baseline	3.885	–	–	–
Semantic Fusion	**1.111**	**0.0039**	**1.00**	**1.00**

Table 7: Main metrics on the SST-2 dataset (validation).

As shown in Table 7, Semantic Fusion reduces the validation PPL from 3.885 (Baseline) to 1.111 (Fusion), $\approx 71.4\%$ relative decrease. The auxiliary semantic MSE = 0.0039, confirming that the model learns the intended fuzzy features.

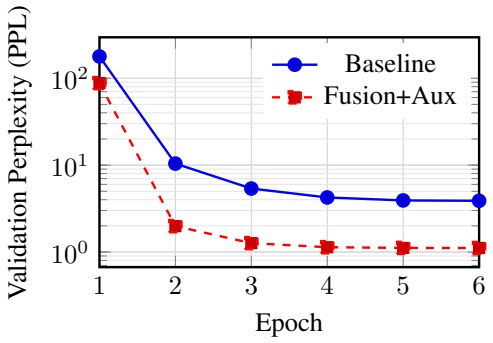

Figure 3: SST-2 validation perplexity across epochs (log scale). Fusion+Aux converges faster and remains better than the baseline.

4 Discussion

Semantic fusion improves language modeling accuracy by injecting an explicit, interpretable semantic channel alongside the text stream. The sentence-level matrix S provides structure at positions with salient choices for this task (e.g. INTENS, ADJ, PUNCT), allowing the backbone to condition on cues it would otherwise have to infer implicitly. A lightweight gated adapter learns to weight this signal where helpful and ignore it where not, while the auxiliary head (with MSE $= 0.0087$) encourages hidden states to align with S. Empirically, this corresponds to lower validation perplexity overall and on the *seen-only* subset, and to reduced token-wise cross-entropy on intensity and punctuation tokens, indicating that gains are not confined to directly supervised adjectives but reflect a broader inductive bias.

The same semantic predicates that the model learns to reconstruct at training time are reused at test time to steer decoding, which makes control robust and transparent. A finite-state *grammar mask* enforces the one-clause template; *state-aware* logit steering targets the adjective and punctuation states; and fuzzy triplets (e.g. `str_low/med/high`) expose graded knobs rather than brittle one-hot tags. Under *hard* constraints (class restriction and deterministic punctuation), polarity and punctuation control are perfect (see Table.2; Table.3 being strictly diagonal). Without hard restriction, *soft* steering is qualitatively effective while preserving diversity (Fig.2).

Beyond accuracy and control, the method generalizes to OOD adjectives by promoting *class semantics* rather than memorized lexemes. A small adjective-class uniformizer discourages within-class mode collapse, and group-uniform mixture sampling (with nucleus applied *after* mixing) preserves probability mass for rare/held-out items; together they yield OOD hit rates of 62% for positive and 43% for negative adjectives (Table.4). The positive rate is on par with a class-uniform reference (3 of 5 held-out items), while the lower negative rate likely reflects milder mixture/truncation settings and the use of medium strength (`str_med`). Practically, the approach is simple and compatible with standard encoder-only LMs (weight tying retained by injecting semantics pre-encoder) and common optimization (*AdamW*, label smoothing, cosine schedule), and it remains interpretable: the feature bank exposes human-readable controls and the auxiliary head's low MSE (0.0087) quantifies how well those semantics are encoded.

5 Conclusion, Limitations and Future Work

Conclusion. We introduced *semantic fusion*, a lightweight mechanism that augments a Transformer LM with a parallel, fuzzy-membership feature channel fused by a gated adapter and regularized with an auxiliary reconstruction head plus a small adjective-class uniformizer. On a controlled synthetic task, it consistently lowers validation perplexity, delivers perfect polarity and punctuation control under hard constraints, generalizes control to held-out adjectives at substantial rates, and yields remarkable reductions in token-level cross-entropy for intensity and punctuation. These gains come with minimal architectural overhead, preserve tied input-output embeddings, and align training-time supervision with test-time controls, providing an interpretable pathway for conditioned natural language generation.

Limitations. Our evaluation uses a synthetic, templated corpus with a small vocabulary and a one-clause decoding grammar (for clarity), which limits ecological validity. The semantic features are hand-specified and depend on fuzzy heuristics (e.g. polarity/strength maps and fixed membership kernels) and our strongest controllability results use *hard* class restriction at decoding; purely *soft* steering on natural text is more challenging. We also report automatic metrics only; human evaluations of fluency, faithfulness, and perceived control are absent. Finally, while training data

often contain two clauses, our qualitative decoding demonstrations predominantly use one-clause FSG, and we did not explore prefix-conditioned prompting in the main experiments.

Future work. To scale beyond synthetic data, we plan to (i) induce the semantic channel from raw text via weak supervision (e.g. lightweight taggers, distant lexicons) or an end-to-end auxiliary predictor trained jointly; (ii) expand the feature bank to tense/aspect, modality, discourse markers, factuality/hedging, coreference, and pragmatic cues; and (iii) replace fixed memberships with learnable, monotonic kernels (e.g. temperature-annealed power laws or spline-based functions) with calibration. For controllability, we aim to strengthen *soft-only* control using contrastive or energy-based heads (or classifier-free-style logit fusion) at designated states, infer grammar state from prefixes, and move beyond a hand-crafted FSG to learned multi-clause constraints. Broader evaluation will include human studies, robustness/safety analyses (e.g. bias-aware features and guardrails), systematic ablations of the gate/auxiliary/uniformizer components, and comparisons against controllable-LM baselines (e.g. CTRL Keskar et al. (2019), PPLM Dathathri et al. (2020), GeDi Krause et al. (2021), DExperts Liu et al. (2021), etc). Finally, we will transfer the approach to open-domain corpora and larger backbones (including causal decoders), explore multilingual settings, and integrate with parameter-efficient adapters, where semantic fusion could complement prompt-based and RLHF-style conditioning.

6 Related work

Transformer LMs and training practices. Modern language modeling is dominated by the *Transformer* architecture Vaswani et al. (2023), whose scalability and capacity to learn from data with minimal inductive bias have enabled strong generative performance across domains such as machine translation Wang et al. (2019b), abstractive summarization Pilault et al. (2020); Kumar & Solanki (2023), open-domain dialogue Roller et al. (2020), and code generation Chen et al. (2021), etc. Several training practices, as adopted in our work, are standard: tying the input and output embeddings reduces parameters and often improves perplexity Shehper et al. (2025); Press & Wolf (2017); Inan et al. (2017); label smoothing mitigates overconfidence and improves generalization Szegedy et al. (2015); Pereyra et al. (2017); *AdamW* Loshchilov & Hutter (2019) decouples weight decay from the adaptive update and is now a default optimizer for LMs Guan (2023); warmup schedules stabilize early training for Transformers Vaswani et al. (2023); and cosine annealing (with or without restarts) is a widely used schedule for large-scale models Loshchilov & Hutter (2017). In our work, we integrate our semantic channel *before* the encoder and keep the hidden width unchanged, so weight tying remains valid by construction Press & Wolf (2017); Inan et al. (2017).

Controllable generation and decoding. Control over style or attributes has been pursued via control codes (conditioning on special tokens) Keskar et al. (2019), plug-and-play guidance with external discriminators that steer logits at inference Dathathri et al. (2020), and classifier-guided priors such as GeDi Krause et al. (2021). Orthogonal to these, constrained decoding restricts the hypothesis space using lexical or grammar constraints Hokamp & Liu (2017). To maintain fluency and reduce degeneration, sampling strategies such as top-k Fan et al. (2018) and nucleus (top-p) Holtzman et al. (2020) sampling are commonly used Holtzman et al. (2020), along with repetition penalties and coverage-style heuristics Holtzman et al. (2020). Our approach differs from prior control methods by (i) exposing an *interpretable* feature channel (polarity, strength, roles) that the model learns to predict and (ii) reusing those same features for test-time steering, while optionally combining with finite-state grammar masks and nucleus sampling for reliability.

Fuzzy logic and auxiliary supervision. Fuzzy sets provide graded membership functions for linguistic categories Zadeh (1965), with standard triangular/Gaussian/trapezoidal parametrizations widely used in fuzzy systems Klir & Yuan (1994). We draw on this formalism to encode token-level predicates as differentiable membership degrees, supplying a low-variance supervisory signal aligned with human-interpretable semantics. Using auxiliary objectives to shape internal representations traces back to multitask learning Caruana (1997); in this work, our auxiliary head predicts the fuzzy feature bank, encouraging hidden states to retain interpretable cues that can be leveraged both for language modeling and for controllable decoding. Finally, our adjective-class "uniformizer" connects to confidence-penalty/entropy-style regularization that discourages overconfident, peaky distributions Pereyra et al. (2017), but is applied *class-conditionally* to promote coverage within a semantic class, aiding OOD control.

References

Rich Caruana. Multitask learning. *Mach. Learn.*, 28(1):41–75, July 1997. ISSN 0885-6125. doi: 10.1023/A:1007379606734. URL https://doi.org/10.1023/A:1007379606734.

Mark Chen, Jerry Tworek, Heewoo Jun, Qiming Yuan, Henrique Ponde de Oliveira Pinto, Jared Kaplan, Harri Edwards, Yuri Burda, Nicholas Joseph, Greg Brockman, Alex Ray, Raul Puri, Gretchen Krueger, Michael Petrov, Heidy Khlaaf, Girish Sastry, Pamela Mishkin, Brooke Chan, Scott Gray, Nick Ryder, Mikhail Pavlov, Alethea Power, Lukasz Kaiser, Mohammad Bavarian, Clemens Winter, Philippe Tillet, Felipe Petroski Such, Dave Cummings, Matthias Plappert, Fotios Chantzis, Elizabeth Barnes, Ariel Herbert-Voss, William Hebgen Guss, Alex Nichol, Alex Paino, Nikolas Tezak, Jie Tang, Igor Babuschkin, Suchir Balaji, Shantanu Jain, William Saunders, Christopher Hesse, Andrew N. Carr, Jan Leike, Josh Achiam, Vedant Misra, Evan Morikawa, Alec Radford, Matthew Knight, Miles Brundage, Mira Murati, Katie Mayer, Peter Welinder, Bob McGrew, Dario Amodei, Sam McCandlish, Ilya Sutskever, and Wojciech Zaremba. Evaluating large language models trained on code, 2021. URL https://arxiv.org/abs/2107.03374.

Sumanth Dathathri, Andrea Madotto, Janice Lan, Jane Hung, Eric Frank, Piero Molino, Jason Yosinski, and Rosanne Liu. Plug and play language models: A simple approach to controlled text generation, 2020. URL https://arxiv.org/abs/1912.02164.

Kawin Ethayarajh. How contextual are contextualized word representations? comparing the geometry of bert, elmo, and gpt-2 embeddings, 2019. URL https://arxiv.org/abs/1909.00512.

Angela Fan, Mike Lewis, and Yann Dauphin. Hierarchical neural story generation, 2018. URL https://arxiv.org/abs/1805.04833.

Lei Guan. Weight prediction boosts the convergence of adamw, 2023. URL https://arxiv.org/abs/2302.00195.

John Hewitt and Christopher D. Manning. A structural probe for finding syntax in word representations. In Jill Burstein, Christy Doran, and Thamar Solorio (eds.), *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pp. 4129–4138, Minneapolis, Minnesota, June 2019. Association for Computational Linguistics. doi: 10.18653/v1/N19-1419. URL https://aclanthology.org/N19-1419/.

Chris Hokamp and Qun Liu. Lexically constrained decoding for sequence generation using grid beam search, 2017. URL https://arxiv.org/abs/1704.07138.

Ari Holtzman, Jan Buys, Li Du, Maxwell Forbes, and Yejin Choi. The curious case of neural text degeneration, 2020. URL https://arxiv.org/abs/1904.09751.

Hakan Inan, Khashayar Khosravi, and Richard Socher. Tying word vectors and word classifiers: A loss framework for language modeling, 2017. URL https://arxiv.org/abs/1611.01462.

Nitish Shirish Keskar, Bryan McCann, Lav R. Varshney, Caiming Xiong, and Richard Socher. Ctrl: A conditional transformer language model for controllable generation, 2019. URL https://arxiv.org/abs/1909.05858.

George J. Klir and Bo Yuan. *Fuzzy sets and fuzzy logic: theory and applications*. Prentice-Hall, Inc., USA, 1994. ISBN 0131011715.

Ben Krause, Akhilesh Deepak Gotmare, Bryan McCann, Nitish Shirish Keskar, Shafiq Joty, Richard Socher, and Nazneen Fatema Rajani. Gedi: Generative discriminator guided sequence generation. In *Findings of Empirical Methods in Natural Language Processing (Findings of EMNLP)*, 2021.

Sandeep Kumar and Anil Solanki. An abstractive text summarization technique using transformer model with self-attention mechanism. *Neural Computing and Applications*, 35:18603–18622, 2023. doi: 10.1007/s00521-023-08687-7. URL https://doi.org/10.1007/s00521-023-08687-7. Published: June 1, 2023; Issue Date: September 2023.

Alisa Liu, Maarten Sap, Ximing Lu, Swabha Swayamdipta, Chandra Bhagavatula, Noah A. Smith, and Yejin Choi. Dexperts: Decoding-time controlled text generation with experts and anti-experts, 2021. URL https://arxiv.org/abs/2105.03023.

Ilya Loshchilov and Frank Hutter. Sgdr: Stochastic gradient descent with warm restarts, 2017. URL https://arxiv.org/abs/1608.03983.

Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. In *International Conference on Learning Representations (ICLR)*, 2019. arXiv:1711.05101.

Tomas Mikolov, Kai Chen, Greg Corrado, and Jeffrey Dean. Efficient estimation of word representations in vector space, 2013. URL https://arxiv.org/abs/1301.3781.

Gabriel Pereyra, George Tucker, Jan Chorowski, Łukasz Kaiser, and Geoffrey Hinton. Regularizing neural networks by penalizing confident output distributions, 2017. URL https://arxiv.org/abs/1701.06548.

Matthew E. Peters, Mark Neumann, Mohit Iyyer, Matt Gardner, Christopher Clark, Kenton Lee, and Luke Zettlemoyer. Deep contextualized word representations, 2018. URL https://arxiv.org/abs/1802.05365.

Jonathan Pilault, Raymond Li, Sandeep Subramanian, and Chris Pal. On extractive and abstractive neural document summarization with transformer language models. In Bonnie Webber, Trevor Cohn, Yulan He, and Yang Liu (eds.), *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pp. 9308–9319, Online, November 2020. Association for Computational Linguistics. doi: 10.18653/v1/2020.emnlp-main.748. URL https://aclanthology.org/2020.emnlp-main.748/.

Ofir Press and Lior Wolf. Using the output embedding to improve language models, 2017. URL https://arxiv.org/abs/1608.05859.

Stephen Roller, Emily Dinan, Naman Goyal, Da Ju, Mary Williamson, Yinhan Liu, Jing Xu, Myle Ott, Kurt Shuster, Eric M. Smith, Y-Lan Boureau, and Jason Weston. Recipes for building an open-domain chatbot, 2020. URL https://arxiv.org/abs/2004.13637.

Ali Shehper, Anibal M. Medina-Mardones, Lucas Fagan, Bartłomiej Lewandowski, Angus Gruen, Yang Qiu, Piotr Kucharski, Zhenghan Wang, and Sergei Gukov. What makes math problems hard for reinforcement learning: a case study, 2025. URL https://arxiv.org/abs/2408.15332.

Richard Socher, Alex Perelygin, Jean Y. Wu, Jason Chuang, Christopher D. Manning, Andrew Y. Ng, and Christopher Potts. Recursive deep models for semantic compositionality over a sentiment treebank. In *Proceedings of the 2013 Conference on Empirical Methods in Natural Language Processing*, pp. 1631–1642. Association for Computational Linguistics, 2013.

Jianlin Su, Yu Lu, Shengfeng Pan, Ahmed Murtadha, Bo Wen, and Yunfeng Liu. Roformer: Enhanced transformer with rotary position embedding, 2023. URL https://arxiv.org/abs/2104.09864.

Christian Szegedy, Vincent Vanhoucke, Sergey Ioffe, Jonathon Shlens, and Zbigniew Wojna. Rethinking the inception architecture for computer vision, 2015. URL https://arxiv.org/abs/1512.00567.

Ian Tenney, Dipanjan Das, and Ellie Pavlick. Bert rediscovers the classical nlp pipeline, 2019. URL https://arxiv.org/abs/1905.05950.

Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. Attention is all you need, 2023. URL https://arxiv.org/abs/1706.03762.

Alex Wang, Amanpreet Singh, Julian Michael, Felix Hill, Omer Levy, and Samuel R. Bowman. GLUE: A multi-task benchmark and analysis platform for natural language understanding. In *Proceedings of the International Conference on Learning Representations (ICLR)*, 2019a.

Qiang Wang, Bei Li, Tong Xiao, Jingbo Zhu, Changliang Li, Derek F. Wong, and Lidia S. Chao. Learning deep transformer models for machine translation. In Anna Korhonen, David Traum, and Lluís Màrquez (eds.), *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pp. 1810–1822, Florence, Italy, July 2019b. Association for Computational Linguistics. doi: 10.18653/v1/P19-1176. URL https://aclanthology.org/P19-1176/.

L.A. Zadeh. Fuzzy sets. *Information and Control*, 8(3):338–353, 1965. ISSN 0019-9958. doi: https://doi.org/10.1016/S0019-9958(65)90241-X. URL https://www.sciencedirect.com/science/article/pii/S001999586590241X.

A Appendix

B Implementation details

Compute environment. Experiments were run in Google Colab, on a Linux x86_64 virtual machine (KVM) with a single-socket Intel Xeon CPU @ 2.20 GHz (6 physical cores / 12 threads), 179 GB RAM, and ∼253 GB local storage. The host supports AVX2/AVX-512 (incl. VNNI). No discrete GPU was available, so all training/inference used the PyTorch CPU backend.

Backbone and parameterization. Both baseline and fusion models use an encoder-only Transformer with *tied* input/output embeddings [7]. Hidden size $d = 128$, layers $L = 4$, heads $H = 4$, feed-forward width 256, dropout 0.1, sinusoidal positional encodings.

Semantic feature bank. For each token we compute $s_t \in [0,1]^F$ over interpretable predicates: `is_noun`, `is_verb`, `is_adj`, `is_subject`, `is_object`, `is_head`, `is_bos`, `is_eos`, `is_comma`, `is_question`, `pos_low/med/high`, `neg_low/med/high`, `str_low/med/high`, `coref_subject`, `is_capitalized`, `is_pronoun`. Graded predicates use fuzzy/triangular memberships with a power kernel

$$\mu(x; c, \tau) = 0.9^{|x-c|/\tau}, \quad \mathrm{Tri}(x; \{0.2, 0.6, 1.0\}, \tau), \;\; \tau = 0.35$$

Fusion module. Let $e_t = E[x_t]$, $u_t = W_s s_t$, and $g_t = \sigma\big(W_g[e_t; s_t]\big)$. The fused input to the encoder is

$$h_t^{(0)} = e_t + u_t + g_t \odot u_t$$

followed by positional encoding and the Transformer encoder. The LM head is tied: $p(y_t|x_{\le t}) = \mathrm{softmax}(E^\top H_t)$. An auxiliary MLP predicts $\hat{s}_t = \sigma(\mathrm{MLP}(H_t))$.

Optimization and schedule. We train for 6 epochs with *AdamW* Loshchilov & Hutter (2019) (learning rate 3×10^{-4}, weight decay 0.01), batch size 64, gradient clipping at norm 1.0, and a warmup+cosine schedule (10% warmup). Label smoothing $\varepsilon = 0.02$ is applied to the LM loss during training (turned *off* for validation metrics). The auxiliary semantic head uses BCE with weight $\lambda_{\text{aux}} = 0.5$. The adjective-class uniformizer has coefficient $\lambda_{\text{uni}} = 0.01$:

$$\mathcal{L}_{\text{uni}} = \frac{1}{|\mathcal{A}|} \sum_{t \in \mathcal{A}} \mathrm{KL}\Big(p_t^{(\mathcal{C}_t)} \Big\| u_{\mathcal{C}_t}\Big), \qquad u_{\mathcal{C}_t}(w) = \tfrac{1}{|\mathcal{C}_t|}$$

Data generation. Synthetic corpus with probability 0.6 of adding a second clause that corefers via pronoun (`she/he/they`). Vocabulary: subjects {`Alice`, `Bob`, `Carol`, `Dave`, `Eve`}; verbs {`finishes`, `reviews`, `trains`, `starts`, `cooks`}; objects {`task`, `paper`, `model`, `project`, `meal`}; intensifiers {`slightly`, `moderately`, `very`, `extremely`}; adjectives (pos) {`good`, `great`, `excellent`, `pleasant`, `wonderful`} and (neg) {`bad`, `poor`, `terrible`, `unpleasant`, `awful`}. We hold out roughly half of the adjectives (both polarities) for OOD control; e.g. pos hold-out {`wonderful`, `excellent`, `great`}, neg hold-out {`terrible`, `awful`, `unpleasant`}. Train/val sizes: $8,000/1,200$; max length 28 (incl. BOS/EOS).

Prompted decoding. If a user-supplied prefix matches the one-clause FSG (e.g. *Carol starts the model,*), we validate it, set the next grammar state accordingly (here, INTENS), and continue decoding from that state. For fusion, semantic features are recomputed for the prefix before continuation.

Decoding. Finite-state grammar (FSG) for one clause:

$$\text{SUBJ} \to \text{VERB} \to \text{“the”} \to \text{OBJ} \to \text{“,”} \to \text{INTENS} \to \text{ADJ} \to \text{PUNCT}.$$

We apply a grammar mask (allowed tokens only), a last-k repetition penalty ($k = 3$; stronger for the baseline, factor ≈ 2.5, vs. fusion ≈ 1.5), and state-aware logit steering at INTENS/ADJ/PUNCT. Under strong requests we optionally hard-restrict adjectives to the desired polarity and deterministically set punctuation. To encourage OOD coverage at ADJ we sample from a group-uniform mixture and then apply nucleus filtering:

$$p = \mathrm{softmax}(\ell_{\mathcal{C}}/T), \qquad u_{\mathcal{C}}(w) = \tfrac{1}{|\mathcal{C}|}\mathbf{1}[w \in \mathcal{C}], \qquad q = (1-\alpha)\,p + \alpha\, u_{\mathcal{C}}, \;\; \alpha \in [0,1].$$

$$\text{Sample using nucleus (top-}\rho\text{) applied to } q$$

Typical settings: baseline fair decoding uses temperature 0.7, top-$p = 0.9$, top-$k = 20$; fusion neutral uses $T = 0.7$, $p = 0.9$; fusion (positive & strong) uses control {`pos_high`$= 0.95$, `str_high`$= 0.9$} with mixture $T = 1.5$, $p = 1.0$, $\alpha = 0.97$; fusion (negative & question) uses {`neg_high`$= 0.95$, `is_question`$= 1.0$, `str_med`$= 0.6$} with $T = 1.3$, $p = 0.95$, $\alpha = 0.85$ (and at PUNCT choose “?” deterministically).

[7] **Tied input/output embeddings.** The output softmax weights reuse the input embedding matrix: if $E \in \mathbb{R}^{V \times d}$ is the token embedding table, set $W_{\text{out}} = E^\top$ so logits are $z_t = h_t E^\top + b$. This reduces parameters and aligns input/output geometry; it requires hidden size d to equal the embedding dimension.

Metrics. Validation PPL is computed with label smoothing *off*. Seen-only PPL masks positions whose gold token is in the held-out adjective set before averaging cross-entropy (Appendix.C). Semantic MSE averages $(\hat{s}_t - s_t)^2$ over non-pad positions. Focus-token CE is the mean cross-entropy conditioned on positions where a given token is the gold target. Control accuracy is measured over $N = 200$ generations per setting (adjective polarity and punctuation). OOD hit rate is the fraction of controlled generations that realize a *held-out* adjective.

Seen-only PPL computation. Let $\mathcal{H}$ be the set of held-out adjective ids. For each position t, we drop the loss if the gold token $y_t \in \mathcal{H}$ and average cross-entropy over the remaining positions; perplexity is $\exp$ of this average. See also Appendix.C.

Reproducibility & configuration. We fix the random seed for `random`, `numpy`, and `torch`, and set the device to `cuda` if available, else `cpu`. All other architectural choices (tied embeddings; 4-layer encoder, $d = 128$, $H = 4$, FFN width 256) and decoding settings are specified in the corresponding sections above to avoid duplication.

C SEEN-ONLY PERPLEXITY

This appendix formalizes the *seen-only* perplexity reported in our experiments and matches the implementation of `eval_perplexity_seen_only` in the code.

Setup and notation. Let $x_{1:L}$ be a validation sequence with gold next-token targets $y_{1:L}$ over vocabulary $\mathcal{V}$, and let $p_\theta(\cdot \mid x_{\le t})$ denote the model's next-token distribution (evaluated *without* label smoothing). Let $\mathcal{H} \subset \mathcal{V}$ be the set of *held-out* adjective token ids used for OOD testing:

$$\mathcal{H} = \big\{\mathrm{stoi}(w) : w \in \texttt{ADJ_POS_HOLD} \cup \texttt{ADJ_NEG_HOLD}\big\}$$

Define the per-position cross-entropy (natural log)

$$\mathrm{CE}_t = -\log p_\theta\big(y_t \big| x_{\le t}\big)$$

and a binary weight that drops padding and held-out targets:

$$w_t = \mathbf{1}\big[y_t \neq \texttt{<pad>}\big] \cdot \mathbf{1}\big[y_t \notin \mathcal{H}\big]$$

Metric. Aggregate the masked loss and token count

$$L_{\text{seen}} = \sum_{t=1}^{L} w_t \mathrm{CE}_t, \qquad N_{\text{seen}} = \sum_{t=1}^{L} w_t$$

and report

$$\mathrm{PPL}_{\text{seen-only}} = \exp\left(\tfrac{L_{\text{seen}}}{N_{\text{seen}}}\right)$$

Batch form. For a batch with logits $\mathbf{Z} \in \mathbb{R}^{B \times L \times |\mathcal{V}|}$, gold $y \in \mathcal{V}^{B \times L}$, and non-pad mask $m \in \{0,1\}^{B \times L}$, define

$$\kappa_{b,t} = \mathbf{1}[y_{b,t} \notin \mathcal{H}], \quad c_{b,t} = \mathrm{CE}\big(\mathrm{softmax}(\mathbf{Z}_{b,t,:}),\ y_{b,t}\big)$$

Then

$$\bar{c}_{\text{seen}} = \frac{\sum_{b,t} c_{b,t} m_{b,t} \kappa_{b,t}}{\sum_{b,t} m_{b,t} \kappa_{b,t}}, \qquad \mathrm{PPL}_{\text{seen-only}} = \exp(\bar{c}_{\text{seen}})$$

Implementation notes. (i) Label smoothing is *disabled* for this evaluation (true one-hot targets). (ii) If $\sum_{b,t} m_{b,t}\kappa_{b,t} = 0$ (no eligible positions), the implementation returns `NaN`. (iii) Context filtering is *not* applied: held-out adjectives may appear in the *inputs* $x_{\le t}$; we only exclude positions whose *targets* are held-out.

Rationale. This metric isolates predictive quality on the *seen* portion of the vocabulary under realistic contexts, which is why $\mathrm{PPL}_{\text{seen-only}}$ is typically lower than overall PPL in our runs.