

000
001
002
003
004
005
006
007
008
009
010
011
012
013
014
015
016
017
018
019
020
021
022
023

SEMANTIC CALIBRATION IN MEDIA STREAMS

003 **Anonymous authors**

004 Paper under double-blind review

ABSTRACT

009 Current generative models can produce synthetic media that is visually indistin-
010 guishable from real content. As a result, traditional detection methods rely mostly
011 on subtle artifacts introduced during generation. However, we show that such
012 methods could eventually become ineffective. Anticipating this, we suggest that
013 the main risk lies not in whether a media sample is synthetic or real, but in whether
014 its semantic content is *deceiving*, that is, whether it distorts the information distri-
015 bution in a way that misrepresents reality. To capture this, we formally introduce
016 the notion of *deception* in the context of online media streams. Complementing
017 standard detection approaches, we introduce *semantic calibration* to mitigate decep-
018 tion directly by processing semantic content using captioning and large language
019 models, rather than relying on artifacts introduced by generative models. Our
020 method is explainable, transparent, and modality agnostic, providing a rigorous
021 foundation for developing new tools to combat online misinformation. We offer
022 both theoretical justification and empirical evidence for its effectiveness.

1 INTRODUCTION

024 Artificial intelligence has opened a new front in the war against misinformation on online platforms
025 (Day, 2019). Modern generative tools are now capable of producing synthetic media, sometimes called
026 *deepfakes* (Ramesh et al., 2022; Rombach et al., 2022; Saharia et al., 2022), that is indistinguishable
027 from real content to the human eye. As a result, detecting these digital forgeries depends increasingly
028 on heuristics and learned methods (Mirsky & Lee, 2021; Verdoliva, 2020; Rana et al., 2022; Heidari
029 et al., 2024). While it may initially appear that the arms race between generation and detection is
030 balanced (i.e., that advances in generative capabilities can, in principle, be met with corresponding
031 progress in detection techniques), we show in this work that under mild assumptions, any detection
032 methods that rely solely on invisible artifacts (i.e., non-semantic content) could eventually become
033 obsolete, echoing concerns that current detection cues may not remain reliable (Wang et al., 2020;
034 Corvi et al., 2023).

035 While this may seem like a discouraging conclusion, we argue that the ultimate goal is not to simply
036 determine whether a media sample is real or synthetic, but rather to assert if its semantic information
037 can be trusted. Detecting synthetic media serves only as a *proxy* for identifying (and potentially
038 removing) deceptive information, based on the assumption that real media reliably reflects true events,
039 but synthetic media does not necessarily convey false information. For example, refining a real
040 but low-resolution image with superresolution tools produces a synthetic image without altering
041 the original semantics (Ledig et al., 2017; Saharia et al., 2023). Conversely, real content can be
042 overrepresented or taken out of context (so-called *cheapfake* (Paris & Donovan, 2019)), posing risks
043 similar to those of synthetic media. Additionally, the boundary between real and generated media is
044 often blurred. For instance, studies have shown that diffusion models can memorize and generate
045 images from their training sets (Somepalli et al., 2023a;b). If a generated image is a pixel-perfect
046 replica of a real photo, should it still be considered fake?

047 Motivated by these observations, we present in this work an alternative, holistic perspective on the
048 issue, by shifting the focus from the *provenance* of a media sample (was it synthetically generated?)
049 to its *semantic plausibility* (is it surprising beyond credible bounds?). We refer to this property as
050 *deception*, capturing the idea that deepfakes and cheapfakes can distort the distribution of semantic
051 information across media streams, and therefore deceive. Building on this perspective, we introduce
052 *semantic calibration*, a method designed to mitigate this risk, even in a future where generative
053 models are no longer detectable. Our approach first converts the semantic content of a media sample

054
055
056
057
058
059

This image shows a male tennis player in action on a clay court. He is wearing a blue shirt, white shorts, and white sneakers. The player is holding a red and black tennis racket and is in the middle of a forehand swing, with his body stretched out to hit the ball. His shadow can be seen on the ground, indicating that he is in motion. The court appears to be well-maintained and the clay is a reddish-brown color.

060
061
062
063
064

This image shows two young men sitting on a couch in a living room. They are both wearing white shirts and black ties. The man on the left is wearing glasses and is smiling at the camera. He is holding a white mug in his hand and appears to be giving it to the other man. There is a basket of flowers on the couch next to them. In the background, there is a piano and books with various items on it.

065
066
067
068
069
070
071
072
073
074

Figure 1: Saliency map showcasing the explainability of semantic calibration on two test images from the COCO dataset. We simulate a setting where sport-related activities are considered fake, i.e. $q(z = \text{sport}) > p_r(z = \text{sport}) = 0$, by creating a dataset \mathcal{D}_r with no sport-related labels. Tokens highlighted in **blue favor acceptance** ($\Delta_i > 0$) while those highlighted in **orange favor rejection** ($\Delta_i < 0$). For clarity, tokens belonging to the same word are merged. As expected, the algorithm rejects the first image and accepts the second with probability ≈ 1 . More importantly, it bases its decisions on words that are intuitive for the given setting (e.g., tennis, forehand, stretched for rejection, and sitting, couch, glasses for acceptance). For the same two images, we also plot the rolling acceptance probabilities in Fig. 5 for three milder semantic shifts $q(z = \text{sport}) > p_r(z = \text{sport}) > 0$.

075
076

into text using a captioning model, then applies rejection sampling in this textual space. Acceptance probabilities are computed using two fine-tuned LLMs trained to model the semantic distributions.

077
078
079
080
081
082
083
084
085
086
087
088
089

Although the distinction between semantic and non-semantic content may seem arbitrary, it reflects how modern generative models are trained (Ramesh et al., 2022; Rombach et al., 2022; Saharia et al., 2022). Semantic content can be understood as the conditioning signal (i.e., prompts), while non-semantic content includes the low-level details the model autonomously fills in to generate realistic outputs. In high-dimensional modalities such as images or video, non-semantic components such as grass patterns, cloud shapes, or water textures dominate the overall data. We estimate in Appendix A.1 that for a small natural image, only about 3% of the information perceived by a human is semantic. Due to this imbalance, traditional deepfake detection methods have remained largely content-agnostic, relying instead on non-semantic artifacts for detection (Frank et al., 2020; Durall et al., 2020; Rössler et al., 2019). However, such approaches fundamentally rely on the imperfections of current generative models, which future models may no longer exhibit.

090
091
092
093
094
095

For these reasons, it is important to explore alternative strategies. We propose reframing deepfakes not as a binary issue of authenticity, but as a distributional problem of semantic information, where the primary risk lies in how a media sample shifts this distribution, regardless of whether it is real or synthetic. To the best of our knowledge, this is the first work to formalize deception in distributional terms and to present an explainable approach for mitigating it. We summarize our contributions as follows:

096
097
098
099
100
101
102
103

- First, under mild assumptions and drawing on classical results from hypothesis testing, we **show that traditional deepfake detection methods could eventually become obsolete** if they rely purely on non-semantic artifacts.
- Anticipating this scenario, we **adopt a holistic, distributional perspective and introduce the concept of deception**: the strategic manipulation of the semantic distribution of media streams. Our analysis shows that deepfake detection serves only as a proxy for our primary goal of reducing semantic deception.
- Finally, we **introduce a method that targets deception directly** by filtering media based on their semantic content. Our approach converts media into text via a captioning model, then applies rejection sampling using two fine-tuned LLMs modeling the semantic distributions. We further provide theoretical evidence that our formulation is well-posed, validate our methods empirically, and explore the explainability of our approach.

108 **2 BACKGROUND**
 109

110 **Media stream.** We consider a setting where various forms of media (e.g., images, videos, audio)
 111 are uploaded to an online platform and redistributed to users of that platform. The core of our analysis
 112 is built on a set of random variables that model the media stream. Specifically, $X \in \mathcal{X}$ denotes the
 113 media itself. Each media item carries semantic content represented by a latent variable $Z \in \mathcal{Z}$ that is
 114 deterministically determined by X , i.e., $Z = f(X)$. To simplify our analysis, we consider that both
 115 \mathcal{X} and \mathcal{Z} are discrete sets (e.g., digital content). We introduce a binary variable $G \in \{0, 1\}$ indicating
 116 whether the media is real ($G = 0$) or generated ($G = 1$). From these random variables, we define the
 117 *media stream* p as the joint distribution

$$118 \quad p(g, z, x) = \mathbb{P}[G = g, Z = z, X = x].$$

119 When the context is clear, we also use p to refer to the marginal distribution over any subset of
 120 these variables (e.g., $p(x) = \mathbb{P}[X = x]$), and we denote by $p_G = \mathbb{P}[G = 1]$ the prior probability of
 121 observing a generated sample in the media stream p . We assume $0 < p_G < 1$. Additionally, we
 122 denote by $q(x, z) = \mathbb{P}[X = x, Z = z | G = 1]$ the distribution of generated media samples and by
 123 $p_r(x, z) = \mathbb{P}[X = x, Z = z | G = 0]$ the distribution of real media samples. With this notation, the
 124 overall media stream distribution p can be written as the mixture

$$125 \quad p(z, x) = p_G q(z, x) + (1 - p_G) p_r(z, x). \quad (1)$$

126 See Appendix A.2 for a detailed summary of the notation used in this work.
 127

128 **Generative models.** Let $q(x|z) = \mathbb{P}[X = x | Z = z, G = 1]$ denote the conditional generative
 129 model used to generate synthetic media. This may correspond to a neural generative model or an
 130 ensemble of such models. These generators are typically trained to minimize the following quantity
 131 (or a relaxed proxy):

$$132 \quad \mathcal{L}(q, p_r) \triangleq \sup_{z \in \mathcal{Z}} D_{\text{KL}}(q(x|z) \| p_r(x|z)), \quad (2)$$

133 where $p_r(x|z)$ denotes the distribution of real media that conveys semantic information z . To
 134 ensure that $p_r(x|z)$ and $q(x|z)$ are well-defined, we assume that $p_r(z), q(z) > 0$ for any z . See
 135 Appendix A.3 for a discussion about the alternative (forward) objective $\mathcal{L}(p_r, q)$.
 136

137 **Deceptive vs neutral media stream.** We say that a media stream p is *deceptive* if $p(z) \neq p_r(z)$.
 138 This directly implies the following factorization (illustrated in Fig. 2):

$$139 \quad p(z, x) = p_G q(x|z) q(z) + (1 - p_G) p_r(x|z) p_r(z).$$

140 Conversely, if no dependency is assumed between the generation label G and the semantic content
 141 Z , such that there is no edge between G and Z in Fig. 2, we refer to the resulting media stream as
 142 *neutral*. To quantify the deviation from neutrality, we define the *deception* of a media stream p as

$$143 \quad \delta(p | p_r) \triangleq D_{\text{KL}}(p(z) \| p_r(z)). \quad (3)$$

145 **Reducing deception.** One objective of content moderation is to reduce the deception of a media
 146 stream p using a decision rule $\phi(x) = \mathbb{P}[F_\phi = 1 | X = x]$, where $F_\phi = 1$ means that the content is
 147 flagged as deceptive. Moderation then happens by removing flagged content, which results in the
 148 filtered media stream

$$149 \quad p^\phi(g, z, x) = \mathbb{P}[G = g, Z = z, X = x | F_\phi = 0], \quad (4)$$

150 Using the convexity of the KL divergence (Cover & Thomas, 2006), we show in Appendix A.4.1 that

$$151 \quad \underbrace{\delta(p^\phi | p_r)}_{\text{direct objective (ours)}} \leq \underbrace{p_G^\phi \delta(q^\phi | p_r) + (1 - p_G^\phi) \delta(p_r^\phi | p_r)}_{\text{proxy objective (deepfake detection)}}, \quad (5)$$

154 where the superscript ϕ indicates conditioning on $F_\phi = 0$ (e.g., $p_G^\phi = \mathbb{P}[G = 1 | F_\phi = 0]$, see
 155 Appendix A.2 for more details on the notation). From Eq. (5), one can identify two avenues to
 156 decrease the deception of p^ϕ . The first is to build a decision rule that minimizes the right-hand side
 157 of the inequality by reducing p_G^ϕ while keeping $p_r^\phi(z)$ close to $p_r(z)$ such that $\delta(p_r^\phi | p_r) \approx 0$. This
 158 is precisely the aim of traditional deepfake detectors that rely on non-semantical cues. However, as
 159 we show in the next section, this paradigm could become ineffective if generative models continue
 160 to improve. In addition, minimizing the right-hand side of Eq. (5) is only a proxy to the primary
 161 objective, which is to minimize $\delta(p^\phi | p_r)$ directly. Therefore, we propose exploring a more direct
 162 alternative in this paper, which is to bring $p^\phi(z)$ closer to $p_r(z)$ without explicitly targeting p_G^ϕ .

162
 163 **Limits of deepfake detection.** We frame deepfake detection as a binary hypothesis testing problem
 164 where the classification error is captured by the indicator variable $E_\phi = \mathbb{1}\{F_\phi \neq G\}$. Given a media
 165 stream p , the objective of deepfake detection is to find a decision rule ϕ that maximizes the expected
 166 accuracy $\mathbb{P}[E_\phi = 0]$. See Fig. 2 for a detailed view of how deepfake detection integrates with the
 167 media stream p . As conditional generative models continue to improve their ability to replicate
 168 reality, it is reasonable to assume that, for any $\epsilon > 0$, a generator q will eventually exist such that
 169 $\mathcal{L}(q, p_r) \leq \epsilon$. Based on this assumption, we derive a fundamental lower bound on the maximal
 170 achievable accuracy of any decision rule.

171 **Theorem 1.** Let $q(z, x) = q(x | z)q(z)$ be such that $\mathcal{L}(q, p_r) \leq \epsilon$. Then

$$\sup_{\phi} \mathbb{P}[E_\phi = 0] \leq \max\{p_G, 1 - p_G\} \left(1 + \sqrt{\frac{\epsilon + \delta(q | p_r)}{2}}\right). \quad (6)$$

172 See Appendix A.4.3 for the full proof. This bound highlights that as generative models improve (i.e.,
 173 $\epsilon \rightarrow 0$), the accuracy of any detector becomes primarily limited by the deception of the media stream.
 174 Consequently, it will become increasingly difficult to construct semantic-agnostic decision rules ϕ
 175 that perform well across arbitrary semantic distributions $q(z)$ (i.e., such that $p_G^\phi \ll p_G$).
 176

3 METHOD

177 **Calibration.** Our goal is to design a decision rule ϕ that
 178 filters a deceptive media stream p into a neutral stream p^ϕ ,
 179 as per Eq. (4), such that $\delta(p^\phi | p_r) \approx 0$. We refer to this
 180 process as *semantic calibration*. The main challenge is
 181 that we do not sample directly from p_r , but from p . How-
 182 ever, for samples $x \sim p$, we assume access to estimates
 183 of both $p(z)$ and $p_r(z)$, where $z = f(x)$ denotes the
 184 semantic content of x (see below for details on how $f(x)$,
 185 $p(z)$, and $p_r(z)$ are approximated in practice). When a
 186 media x with semantics $z = f(x)$ is sampled from the full
 187 stream p , we accept it with probability $p_r(z)/M p(z)$, where
 188 $M \geq M^* \triangleq \sup_{z' \in \mathcal{Z}} p_r(z')/p(z')$. Using this rejection sam-
 189 pling procedure, it can be shown that $p^\phi(z) = p_r(z)$ (see
 190 Appendix A.4.4).
 191

192 **Computing M^* .** The constant M^* is well-defined only if the support of $p(z)$ is contained within
 193 the support of $p_r(z)$. Moreover, computing M^* directly appears intractable, as it would require
 194 evaluating the ratio $p_r(z)/p(z)$ for every possible semantic representation z . Fortunately, in our setting,
 195 $p(z)$ is a mixture of $q(z)$ and $p_r(z)$ as per Eq. (1). This implies that for any z such that $p_r(z) > 0$,
 196 we also have $p(z) > 0$ (assuming $p_G < 1$). In fact, it can be shown that $M^* \leq 1/(1-p_G)$ (see
 197 Appendix A.4.5). Hence, choosing $M \geq 1/(1-p_G)$ ensures that $p^\phi(z) = p_r(z)$. If p_G is unknown, a
 198 conservative estimate $\hat{p}_G \geq p_G$ can be used instead.
 199

200 **Estimating $p_r(z)/p(z)$.** While X is observed, Z is latent and it remains unclear how to compute the
 201 ratio $p_r(z)/p(z)$ for a given $(x, z) \sim p$. Since z is the semantical information of a media, we propose
 202 to approximate it using a textual representation (i.e. a finite series of n tokens $\hat{Z} = \hat{Z}_1 \hat{Z}_2 \dots \hat{Z}_n$).
 203 To that end, we use a captioning model \hat{f} to estimate f . Ideally, $\hat{z} = \hat{f}(x)$ captures all semantic infor-
 204 mation a human would perceive in media x (formally, we aim for $H(Z | \hat{Z}) = 0$, where $H(Z | \hat{Z})$
 205 denotes the entropy of Z given \hat{Z}). To achieve this, the model should be biased toward generating
 206 detailed captions. This reduces the risk of omitting semantically relevant details. However, because
 207 the mapping $z \mapsto \hat{z}$ is one-to-many, the ratios $p_r(\hat{z})/p(\hat{z})$ might differ from $p_r(z)/p(z)$. Fortunately,
 208 under the assumption of perfect generation, i.e., $q(x|z) = p_r(x|z)$ for any z , we show that both ratios
 209 are equal (see Appendix A.4.6).
 210

211 **Estimating $p_r(\hat{z})/p(\hat{z})$.** Once $\hat{z} = \hat{f}(x)$ is computed, we estimate $p_r(\hat{z})$ and $p(\hat{z})$ using language
 212 models $\pi_{\theta_r}(\hat{z})$ and $\pi_\theta(\hat{z})$ that were auto-regressively trained on the captioned data from the real and
 213 full media streams (\mathcal{D}_r , respectively \mathcal{D}). See Section 6 for more details on how \mathcal{D} and \mathcal{D}_r can be

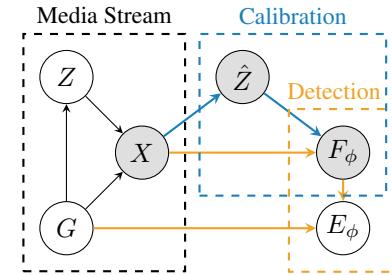


Figure 2: Bayesian network showing the factorization of both **deepfake detection** and **semantic calibration**. Observed variables are shown in gray.

constructed in practice. Finally, we perform rejection sampling in the space of textual approximations \hat{z} , using the ratio $\pi_{\theta_r}(\hat{z})/\pi_{\theta}(\hat{z})$ to approximate $p_r(z)/p(z)$. For numerical stability, we compute this ratio in log-space and only consider a subset of tokens $\mathcal{I}_{\rho}(\hat{z}) \subseteq \{1, \dots, |\hat{z}|\}$ determined by a parameter ρ (see Section 4), leading to the following definition:

$$r(\hat{z}; \pi_{\theta}, \pi_{\theta_r}, \rho) \triangleq \exp \left(\sum_{i \in \mathcal{I}_{\rho}(\hat{z})} [\log \pi_{\theta_r}(\hat{z}_i | \hat{z}_{<i}) - \log \pi_{\theta}(\hat{z}_i | \hat{z}_{<i})] \right) \approx \frac{p_r(\hat{z})}{p(\hat{z})}. \quad (7)$$

Proposed method. Given any media sample x , a captioning model \hat{f} , two language models π_{θ} and π_{θ_r} , and a conservative estimate $\hat{p}_G \geq p_G$, we define the decision rule

$$\phi(x; \pi_{\theta}, \pi_{\theta_r}, \hat{f}, \hat{p}_G, \rho) \triangleq \min \left\{ (1 - \hat{p}_G) \cdot r(\hat{f}(x); \pi_{\theta}, \pi_{\theta_r}, \rho), 1 \right\}, \quad (8)$$

where the clipping ensures the output remains a valid probability by accounting for numerical and approximation errors. Then, when a media sample x is uploaded to the media stream, we sample $F_{\phi} \sim \text{Bernoulli}(\phi(x; \pi_{\theta}, \pi_{\theta_r}, \hat{f}, \hat{p}_G))$, and filter out x when $F_{\phi} = 1$.

4 EXPERIMENTS

Overview. We conduct a series of experiments to evaluate the effectiveness of our decision rule for mitigating deception in media streams. To simulate the case where generation is perfect, we rely exclusively on real text, image, and audio datasets, therefore ensuring $\mathcal{L}(q, p_r) = 0$. Firstly, we quantitatively evaluate our semantic calibration method. This allows us to work with known semantic distributions and precisely control distributional shifts. We create these synthetic shifts using labeled datasets and compute deception based on changes in label distributions. Secondly, since real-world scenarios typically lack access to such ground-truth labels, we demonstrate that our method offers high explainability and transparency, making it practical even when deception cannot be directly measured.

Tractable semantic distributions. To create a setting where deception can be evaluated, we use the labeled categorical datasets listed below and assume that the latent semantic variable Z corresponds to the class label. The semantic distribution $p(z)$ is therefore the distribution over these labels, which allows us to create a setting where deception is tractable. We construct two versions of each dataset: one for training the *general* model $\pi_{\theta}(\hat{z})$, denoted by \mathcal{D} , and one for training the *real* model $\pi_{\theta_r}(\hat{z})$, denoted by \mathcal{D}_r . Semantic distortion is introduced by reweighting the label frequencies in \mathcal{D}_r . For each dataset, we set target class proportions and fill the real training set using as much data as possible without exceeding the specified counts. We generate multiple distorted versions per dataset, each corresponding to a different label distribution. The severity of each shift is quantified using $\delta(p | p_r)$.

Datasets. Our experiments span four modalities: text, image, video, and audio. For text, we use Senator Tweets (KristinCoyote, 2025) with conservative/liberal labels, Political Podcasts (nbandhi, 2025) with again the same labeling. We also use Political Bias (surajkarakulath, 2025) with labels left/center/right, AG-News (Zhang et al., 2015), labeled across four categories, and a custom dataset from eight distinct books from Project Gutenberg (Project Gutenberg, 2025), where semantically split phrases are labeled by book origin. For images, we use CIFAR-10 (Krizhevsky, 2009) (10 classes), CIFAR-100 (Krizhevsky, 2009) (100 classes), a coarse-labeled ImageNet (Deng et al., 2009) subset (20 custom coarse classes), and a COCO (Lin et al., 2014a) subset with a binary coarse labeling. We also apply our framework in the multi-modal setting of images with text. For this, we use Hateful

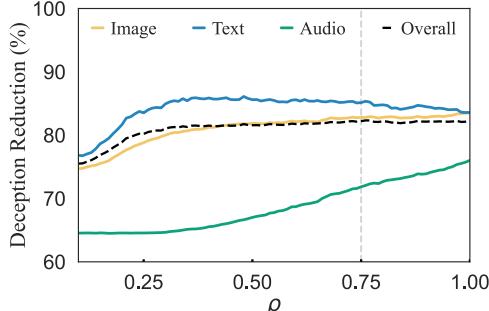


Figure 3: Filtering performance across ρ values. We use $\rho = 0.75$ for text and image, $\rho = 1$ for audio.

270 Memes Challenge Kiela et al. (2021) with existing labels (hateful and non-hateful). For video, we use
 271 ActivityNet (Caba Heilbron et al., 2015), with 52 coarse labels. For audio, we use VGGSound (Chen
 272 et al., 2020) (14 custom coarse classes) and UrbanSound8k (Salamon et al., 2014) (10 original
 273 classes). A detailed description of each dataset and labeling scheme for each experiment can be
 274 found in Appendix A.5. We generate image captions using the Florence (Yuan et al., 2021) model
 275 for all datasets. For audio, we use Qwen-Audio (Chu et al., 2023) to caption UrbanSound8k and
 276 reuse AudioSetCaps (Mei et al., 2023) captions for VGGSound. The exact prompts used to caption
 277 the different modalities can be found in Appendix A.6. Each original dataset is first split 80/20 into
 278 general training and test subsets, and the general training set is then further split 80/20 into training
 279 (\mathcal{D} and \mathcal{D}_r) and validation subsets.

280 **Training.** Our pipeline begins by translating media into text where applicable (e.g., images via the
 281 Florence captioning model (Yuan et al., 2021), audio via Qwen-Audio (Chu et al., 2023)). We then
 282 fine-tune two GPT-2 Small models (Radford et al., 2019) (124M parameters), one on \mathcal{D} and one on
 283 \mathcal{D}_r , both initialized from pre-trained weights. Smaller LLMs are easier to train for our objective,
 284 especially given the limited training data. Both models are trained for 5 epochs using AdamW
 285 (Loshchilov & Hutter, 2019) with a learning rate of $2e-5$, a linear scheduler, and 1000 warmup steps.
 286 The checkpoint with the best validation loss is selected. Full implementation details can be found in
 287 Appendix A.7.

288 **Inference.** During inference, we evaluate the likelihood ratio of each text or caption $\hat{z} = \hat{f}(x)$ from
 289 the general test set using both models $\pi_\theta(\hat{z})$ and $\pi_{\theta_r}(\hat{z})$, following Eq. (7). To reduce noise from
 290 low-impact tokens and focus on the most informative parts of the sequence, we apply a top- ρ filtering
 291 strategy over the log-probability differences, inspired by nucleus sampling (Holtzman et al., 2020).
 292 Formally, let

$$\Delta_i \triangleq \log \pi_{\theta_r}(\hat{z}_i | \hat{z}_{<i}) - \log \pi_\theta(\hat{z}_i | \hat{z}_{<i}), \quad (9)$$

293 and let ω be a permutation of $\{1, 2, \dots, |\hat{z}|\}$ such that $|\Delta_{\omega(1)}| \geq |\Delta_{\omega(2)}| \geq \dots \geq |\Delta_{\omega(|\hat{z}|)}|$.
 294 We define k_ρ as the smallest k such that $\sum_{j=1}^k |\Delta_{\omega(j)}| \geq \rho \cdot \sum_{i=1}^{|\hat{z}|} |\Delta_i|$, and we use
 295 $\mathcal{I}_\rho(\hat{z}) \triangleq \{\omega(j) \mid 1 \leq j \leq k_\rho\}$ in Eq. (7). This limits the computation to tokens that contribute
 296 most to the distributional shift, improving robustness in long captions. See Fig. 3 for an analysis of
 297 the optimal ρ value for each modality. Finally, we simply use our decision rule in Eq. (8) to compute
 298 the probability of flagging the media sample.

301 5 RESULTS

302 **Quantitative.** We present the main results of our experiments in Table 1. For each dataset, we
 303 evaluate performance across four levels of semantic shift (defined as the deception of the unfiltered
 304 general distribution $\delta(p | p_r)$), including a baseline serving as a no-shift control. For the baselines,
 305 we observe that calibration introduces little to no deception (≤ 0.103 , which, for comparison, is less
 306 than half the level observed in the most deceptive mild-shift setting), confirming that the method does
 307 not introduce unwanted semantic bias when none exists. As the semantic shift increases (from mild to
 308 severe), our method consistently reduces deception across all datasets and modalities, with reductions
 309 ranging from 64% to as high as 99%. Calibration is especially effective for long-text modalities (text
 310 and images) and remains robust for more compact representations like audio captions.

311 **Qualitative.** To complement scalar metrics and provide an intuitive understanding of what different
 312 levels of deception represent, we visually compare in Fig. 4 the filtered and real distributions p^ϕ and
 313 p , respectively. These plots show that our method reliably aligns the filtered distributions with the
 314 intended targets, across all modalities and even under strong semantic shifts. We emphasize that
 315 for all results presented in Table 1 and Fig. 4, labels are only used to measure deception reduction.
 316 They are never used during training, reflecting the fact that Z is latent, as illustrated in Fig. 2. This
 317 demonstrates the potential of our approach to filter media streams based on intractable latent semantic
 318 distributions.

319 **Explainability.** The qualitative and quantitative results presented above are made possible by a
 320 simplified, controlled setting where the semantic space Z is small, discrete, and where the
 321 distributions $p_r(z)$ and $p(z)$ are known. In real-world applications, these semantic distributions are

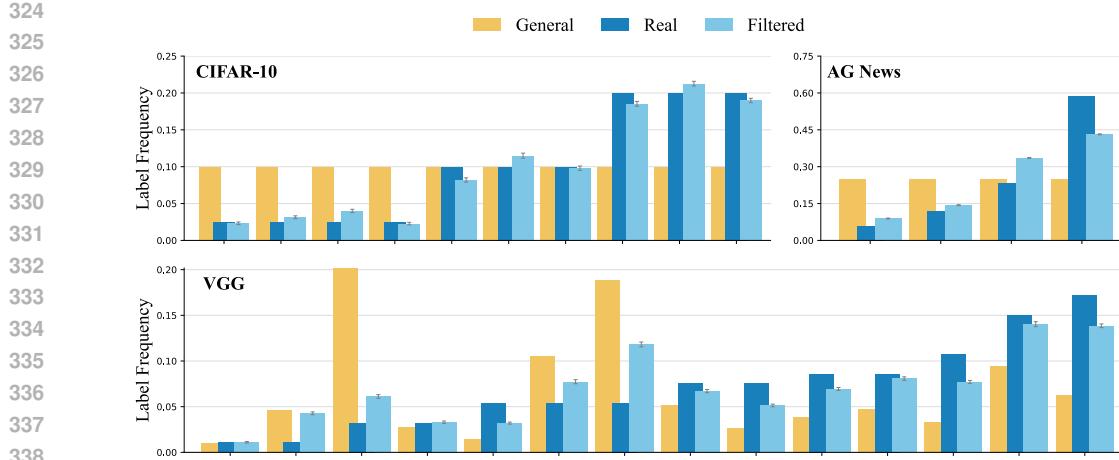


Figure 4: Visualization of the **filtered distribution** $p^\phi(z)$ compared to the **real distribution** $p_r(z)$ and the **general distribution** $p(z)$. The plots show the moderate shift scenario for three datasets CIFAR-10, AG-News, and VGG-Sound, each representing a different modality. The positions along the horizontal axis correspond to class labels, ordered by increasing frequency under $p_r(z)$. We observe that $p^\phi(z)$ generally lies between $p(z)$ and $p_r(z)$, but most closely resembles the latter.

intractable, making direct quantitative evaluation of deception impossible. Nonetheless, our approach provides high explainability and transparency, offering actionable insights even when exact semantic distributions cannot be computed. This constitutes a strong advantage compared to more opaque media stream moderation algorithms. More precisely, we can gain fine-grained insight into the impact of each word by analyzing the per-token log-probability differences Δ_i . This makes it possible to identify which parts of the caption are semantically more typical of the real distribution versus the general one, shedding light on how our filtering algorithm makes its rejection decisions. For instance, Fig. 1 presents a token-level saliency map for two COCO images, highlighting the semantic cues the algorithm relies on to make its decisions. Additional figures can be found in the Appendix A.11. Additionally, Fig. 5 illustrates how the severity of the semantic shift affects rejection behavior, with milder shifts leading to more uncertain acceptance or rejection decisions, as expected.

6 DISCUSSION AND LIMITATIONS

Relation to out-of-distribution (OOD) detection. OOD detection is the binary task of determining whether a sample comes from a reference distribution. When operating in the semantic space \mathcal{Z} , OOD methods typically rely on $p_r(z)$, classifying a media sample x with semantics $z = f(x)$ as OOD whenever $p_r(z) < \tau$ for a given threshold $\tau > 0$ (methods that do not use p_r are quite disconnected from our setting). If x is observed twice, it is expected that the detection mechanism will come up with the same conclusion. Each sample x is treated independently and the actual distribution of x is typically of no concern. Semantic calibration, on the other hand, is interested in the ratio $p_r(z)/p(z)$, with $p(z)$ being a reference distribution. If both $p_r(z)$ and $p(z)$ are small but similar, then calibration might accept the media x , whereas an OOD detector might reject the sample due to its low semantic probability. However, there is a special case in which semantic calibration coincides with OOD detection, namely when p_r and q have disjoint support. In that case, $p_z = (1 - p_G)p_r(z)$ for samples

Table 2: Detection performance of Eq. (8) on misinformation and hate detection datasets: ISOT (Ahmed et al., 2017), LIAR (Wang, 2017), FAKE TFG (Álvarez Hervás, 2022) and Hateful Memes Kiela et al. (2021).

| Dataset | Acc. | Recall | Prec. | F1 |
|---------------|-------|--------|-------|-------|
| ISOT | 0.998 | 0.999 | 0.997 | 0.998 |
| LIAR | 0.869 | 0.851 | 0.952 | 0.887 |
| Fake TFG | 0.972 | 0.977 | 0.972 | 0.974 |
| Hateful Memes | 0.842 | 0.890 | 0.729 | 0.802 |

378
379 Table 1: Semantic calibration performance across various datasets and semantic shifts, mea-
380 sured by the deception reduction $(\delta(p|p_r) - \delta(p^\phi|p_r)) / \delta(p|p_r)$. Errors in *italics* indicate *standard de-*
381 *viation over 1000 inference runs* (for a single training run), while errors in underline indicate
382 standard deviation over 3 training runs (for a single inference run). We report training errors for
383 severe shifts only to save computation, as similar errors were observed across other shifts. We use the
384 conservative estimate $\hat{p}_G = 1 - |\mathcal{D}_r| / |\mathcal{D}|$. The shifts shown in **blue** can be visualized in Fig. 4.

| | Dataset | $ \mathcal{D}_r $ | \hat{p}_G | Shift | $\delta(p p_r)$ | $\delta(p^\phi p_r)$ | Reduction |
|-------------------------------------|----------------------------------|-------------------|-------------|-----------------|-----------------|---------------------------------------|------------------|
| Text $(\rho = 0.75)$ | Senator Tweets | 31.6k | 0.60 | Baseline | 0.000 | 0.000 ± 0.000 | - |
| | $ \mathcal{Z} = 2$ | 35.6k | 0.55 | Mild | 0.020 | 0.001 ± 0.002 | 96.03% |
| | $ \mathcal{D} = 79.2k$ | 33.3k | 0.58 | Moderate | 0.193 | 0.016 ± 0.001 | 91.55% |
| | $\mathbb{E}[\hat{Z}] = 41.3$ | 23.8k | 0.70 | Severe | 0.368 | $0.067 \pm 0.002 / 0.004$ | 81.99% |
| | Political Podcasts | 4.6k | 0.60 | Baseline | 0.000 | 0.001 ± 0.000 | - |
| | $ \mathcal{Z} = 2$ | 9.1k | 0.20 | Mild | 0.022 | 0.000 ± 0.000 | 99.77% |
| | $ \mathcal{D} = 11.39k$ | 7.1k | 0.38 | Moderate | 0.081 | 0.000 ± 0.000 | 99.99% |
| | $\mathbb{E}[\hat{Z}] = 161.0$ | 5.7k | 0.50 | Severe | 0.207 | $0.013 \pm 0.001 / \underline{0.004}$ | 93.79% |
| | Political Bias | 3.8k | 0.60 | Baseline | 0.000 | 0.000 ± 0.000 | - |
| | $ \mathcal{Z} = 3$ | 5.7k | 0.41 | Mild | 0.134 | 0.020 ± 0.003 | 84.85% |
| | $ \mathcal{D} = 9.59k$ | 3.4k | 0.65 | Moderate | 0.297 | 0.001 ± 0.000 | 97.12% |
| | $\mathbb{E}[\hat{Z}] = 1376.9$ | 3.0k | 0.69 | Severe | 0.368 | $0.057 \pm 0.004 / \underline{0.007}$ | 84.52% |
| Image $(\rho = 0.75)$ | AG-News | 38.4k | 0.60 | Baseline | 0.000 | 0.001 ± 0.000 | - |
| | $ \mathcal{Z} = 4$ | 48.0k | 0.50 | Mild | 0.239 | 0.014 ± 0.002 | 94.14% |
| | $ \mathcal{D} = 96.0k$ | 40.8k | 0.58 | Moderate | 0.351 | 0.051 ± 0.001 | 85.47% |
| | $\mathbb{E}[\hat{Z}] = 52.1$ | 34.3k | 0.64 | Severe | 0.603 | $0.006 \pm 0.002 / \underline{0.005}$ | 99.00% |
| | Gutenberg | 1640 | 0.60 | Baseline | 0.000 | 0.103 ± 0.009 | - |
| | $ \mathcal{Z} = 8$ | 2576 | 0.27 | Mild | 0.254 | 0.052 ± 0.005 | 79.52% |
| | $ \mathcal{D} = 4096$ | 1728 | 0.58 | Moderate | 0.563 | 0.071 ± 0.010 | 87.39% |
| | $\mathbb{E}[\hat{Z}] = 109.4$ | 1288 | 0.69 | Severe | 0.860 | $0.014 \pm 0.003 / \underline{0.008}$ | 98.37% |
| | COCO | 17.6k | 0.60 | Baseline | 0.000 | 0.000 ± 0.000 | - |
| | $ \mathcal{Z} = 2$ | 11.5k | 0.74 | Mild | 0.678 | 0.141 ± 0.007 | 79.20% |
| | $ \mathcal{D} = 44.0k$ | 9.5k | 0.78 | Moderate | 1.379 | 0.042 ± 0.004 | 96.95% |
| | $\mathbb{E}[\hat{Z}] = 100.3$ | 9.0k | 0.80 | Severe | 1.926 | $0.009 \pm 0.002 / \underline{0.003}$ | 99.53% |
| Image + Text $(\rho = 1)$ | CIFAR-10 | 16.0k | 0.60 | Baseline | 0.000 | 0.015 ± 0.003 | - |
| | $ \mathcal{Z} = 10$ | 13.4k | 0.67 | Mild | 0.237 | 0.009 ± 0.002 | 96.20% |
| | $ \mathcal{D} = 40.0k$ | 20.0k | 0.50 | Moderate | 0.347 | 0.009 ± 0.002 | 96.95% |
| | $\mathbb{E}[\hat{Z}] = 93.0$ | 11.5k | 0.71 | Severe | 0.781 | $0.059 \pm 0.006 / \underline{0.001}$ | 92.45% |
| | CIFAR-100 | 16.0k | 0.60 | Baseline | 0.000 | 0.073 ± 0.007 | - |
| | $ \mathcal{Z} = 100$ | 14.3k | 0.64 | Mild | 0.216 | 0.076 ± 0.007 | 64.81% |
| | $ \mathcal{D} = 40.0k$ | 14.2k | 0.65 | Moderate | 0.403 | 0.080 ± 0.007 | 80.34% |
| | $\mathbb{E}[\hat{Z}] = 91.6$ | 15.2k | 0.62 | Severe | 0.818 | $0.144 \pm 0.012 / \underline{0.003}$ | 80.15% |
| | ImageNet | 32.5k | 0.60 | Baseline | 0.000 | 0.020 ± 0.003 | - |
| | $ \mathcal{Z} = 20$ | 23.2k | 0.71 | Mild | 0.272 | 0.073 ± 0.003 | 73.16% |
| | $ \mathcal{D} = 81.3k$ | 21.0k | 0.74 | Moderate | 0.526 | 0.106 ± 0.003 | 79.85% |
| | $\mathbb{E}[\hat{Z}] = 98.4$ | 12.6k | 0.85 | Severe | 0.796 | $0.083 \pm 0.006 / \underline{0.003}$ | 89.87% |
| Video $(\rho = 1)$ | Hateful Memes | 2.7k | 0.60 | Baseline | 0.000 | 0.000 ± 0.000 | - |
| | $ \mathcal{Z} = 2$ | 3.4k | 0.50 | Mild | 0.280 | 0.001 ± 0.000 | 99.53% |
| | $ \mathcal{D} = 6.8k$ | 2.9k | 0.57 | Moderate | 0.580 | 0.015 ± 0.000 | 97.40% |
| | $\mathbb{E}[\hat{Z}] = 109.4$ | 2.8k | 0.59 | Severe | 0.781 | $0.043 \pm 0.001 / \underline{0.002}$ | 94.48% |
| | Activity Net | 4.8k | 0.60 | Baseline | 0.000 | 0.000 ± 0.000 | - |
| | $ \mathcal{Z} = 52$ | 9.9k | 0.17 | Mild | 0.295 | 0.024 ± 0.001 | 91.81% |
| | $ \mathcal{D} = 11.9k$ | 7.7k | 0.35 | Moderate | 0.599 | 0.037 ± 0.002 | 93.84% |
| | $\mathbb{E}[\hat{Z}] = 9.8$ | 4.7k | 0.61 | Severe | 0.791 | $0.051 \pm 0.001 / \underline{0.009}$ | 93.55% |
| | Urbansound8k | 1984 | 0.60 | Baseline | 0.000 | 0.011 ± 0.003 | - |
| | $ \mathcal{Z} = 10$ | 2296 | 0.54 | Mild | 0.202 | 0.055 ± 0.006 | 72.77% |
| | $ \mathcal{D} = 4960$ | 2128 | 0.57 | Moderate | 0.616 | 0.188 ± 0.014 | 69.48% |
| | $\mathbb{E}[\hat{Z}] = 15.1$ | 2160 | 0.56 | Severe | 0.825 | $0.221 \pm 0.019 / \underline{0.005}$ | 73.21% |
| Audio $(\rho = 1)$ | VGG-Sound | 58.3k | 0.60 | Baseline | 0.000 | 0.002 ± 0.000 | - |
| | $ \mathcal{Z} = 14$ | 35.5k | 0.76 | Mild | 0.300 | 0.082 ± 0.003 | 72.67% |
| | $ \mathcal{D} = 145.6k$ | 32.0k | 0.78 | Moderate | 0.639 | 0.091 ± 0.003 | 85.76% |
| | $\mathbb{E}[\hat{Z}] = 13.1$ | 29.1k | 0.80 | Severe | 0.927 | $0.152 \pm 0.004 / \underline{0.005}$ | 83.48% |

427
428 in the support of p_r , and $p(z) = p_G q(z)$ otherwise (see Eq. (1)). Setting $M = 1/(1 - p_G)$ as
429 mentioned in Section 3, the ratio $p_r(z)/(Mp(z))$ is either 1 (for z in the support of p_r) or 0 (for z in
430 the support of q). Our filtering rule in Eq. (8) becomes therefore de facto an exact binary, detection

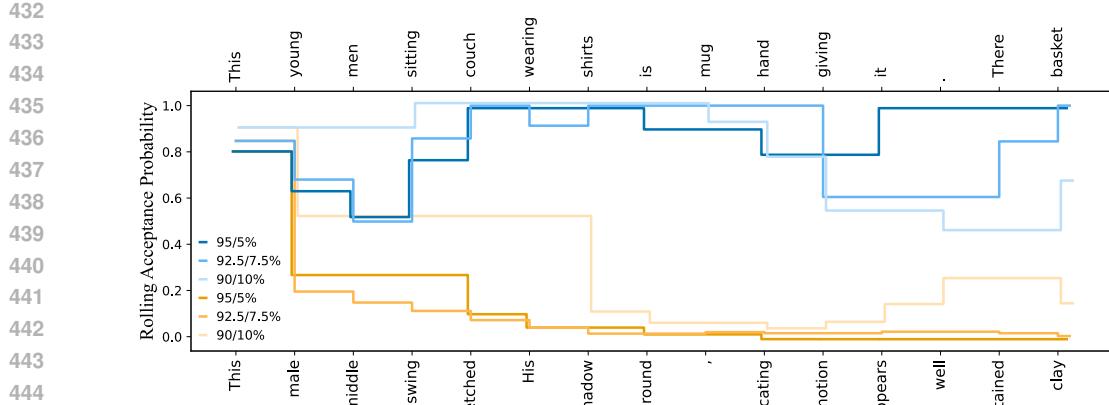


Figure 5: Rolling acceptance probabilities for the two images shown in Fig. 1, in a setting where sport-related content is overrepresented in the general distribution, i.e., $q(z = \text{sport}) > p_r(z = \text{sport})$. Rolling probabilities are computed using Eq. (8) with $\rho = 0.75$, applying a truncated sum in the ratio from Eq. (7). Each image is evaluated under three shift severities in the no-sport/sport scenario. Orange lines show results for the **tennis player**, while blue lines show results for the **two men celebrating**. The general distribution $p(z)$ is 80% no sport, 20% sport.

rule. This special case is precisely the setting of misinformation datasets, as information is either real or fake, but cannot be both (i.e. disjoint support). In Table 2 we present the detection performance of our framework based on our decision rule on a few misinformation and hate detection datasets. The results confirm that our approach generalizes to *semantic detection* tasks on real data.

Semantic calibration is *not* deepfake detection. We emphasize that our approach is not designed to perform deepfake detection. In broad terms, deepfake detection can be seen as an OOD detection task in the media space \mathcal{X} . In contrast, semantic calibration functions exclusively in the semantic latent space \mathcal{Z} , a choice motivated by the fundamental limitations highlighted in Theorem 1. By design, media samples that have the same semantic information are treated equally by semantic calibration, regardless of their truthfulness (generated vs real). Our method is therefore *not* intended as a replacement for existing detection systems, but rather as a complementary framework. An illustrative example highlighting the difference is shown in Appendix A.9.

Semantic calibration is *not* fact-checking. Our approach mitigates deception by measuring deviations from a reference distribution of semantic content, rather than verifying factual correctness. This has several inherent limitations. On one hand, it can flag content that appears abnormally surprising, even if it is factually accurate. Conversely, it may allow content that is factually incorrect to pass undetected if its semantic profile aligns with patterns previously observed in real data. For these reasons, fact-checking should be addressed separately.

Reliance on a trusted dataset \mathcal{D}_r . Another central limitation of our method lies in its reliance on \mathcal{D}_r to approximate the distribution of real content p_r . This assumes that the reference distribution is both representative and sufficiently comprehensive, a strong assumption in dynamic or underrepresented domains. However, this challenge can be mitigated by periodically updating \mathcal{D}_r with new, vetted content, and by retraining or fine-tuning the reference model π_{θ_r} . We refer to p_r as the “real” distribution, but *this is not meant to define what is objectively real*. Instead, p_r should represent a distribution that is broadly regarded as desirable, but we do not address the governance issue on how such a reference distribution is selected. We have primarily framed our analysis around misinformation introduced by deepfakes (hence the term “real”), but semantic calibration is versatile and is not tied to that setting. For instance, three of our experiments address the problem of political balancing, i.e., balancing streams of political tweets, political podcast segments, or political news articles. In all such cases, \mathcal{D}_r is not meant to represent objective reality, it simply specifies the semantic distribution the system is designed to enforce. Without semantic calibration, the semantic distribution of the media stream is more vulnerable to adversarial users.

486 **Reliance on foundation models.** Our method inherits the limitations of current captioning and
 487 language models, which may reduce the reliability of decisions in certain scenarios. Nonetheless, it
 488 employs off-the-shelf captioning models without fine-tuning, demonstrating robustness. Furthermore,
 489 benchmarks such as CapArena Cheng et al. (2025) show that advanced VLMs like GPT-4o can match
 490 or exceed human performance in detailed captioning. Future work should focus on enhancing these
 491 models for semantic calibration.

492 **Broader impact.** The main risk of our method lies in the choice of the reference dataset \mathcal{D}_r , as
 493 any bias it contains will affect the filtered media stream. To mitigate this, we recommend transparent
 494 selection and public release of the datasets \mathcal{D}_r and \mathcal{D} , the captioning model \hat{f} , and the two fine-tuned
 495 LLMs π_θ and π_{θ_r} . If implemented responsibly, our approach could enable content moderation that
 496 is explainable, transparent, and scalable. Additionally, while our results primarily demonstrate the
 497 effectiveness of our approach for filtering media streams, it is unlikely that this would be its first
 498 application in practice. A more practical use would be to apply it as a flagging algorithm in content
 499 moderation, offering moderators (or even other moderation algorithms) a new method to assess the
 500 danger in new media samples.

502 7 RELATED WORK

503 **Fundamental limits of deepfake detection.** While the limitations of deepfake detection are well-
 504 documented (Dolhansky et al., 2020; Wen et al., 2022; Hussain et al., 2021; Ikram et al., 2024), to the
 505 best of our knowledge, very few studies have attempted to establish theoretical bounds on deepfake
 506 detection performance limits in the context of online media streams. Most similar to our approach,
 507 Agarwal & Varshney (2019) frame deepfake detection as a hypothesis testing problem, deriving
 508 performance bounds based on robust statistics. Their analysis, however, is primarily focused on the
 509 context of generative adversarial networks (Goodfellow et al., 2014). Our work extends this line
 510 of reasoning by providing a theoretical bound applicable to a more general setting, independent of
 511 the specific generative model architecture, further motivating the need for alternative approaches as
 512 generative capabilities advance. To the best of our knowledge, we are the first to explicitly distinguish
 513 the semantic and non-semantic information in a media sample in the context of online media forensics.

514 **Content Moderation.** Automated content moderation has evolved from transformer-based classi-
 515 fiers with lightweight toxicity heads (Lees et al., 2022), to multimodal architectures combining text
 516 and image encoders with cross-modal contrastive training (Yuan et al., 2023), and more recently to
 517 instruction-tuned multimodal assistants that directly answer whether content violates guidelines and
 518 generate explanations (Wu et al., 2024). In parallel, LLMs have opened new avenues for addressing
 519 problematic content through semantic understanding, being applied to automated fact-checking
 520 (Vykopal et al., 2024; Kotonya & Toni, 2020; Hu et al., 2024) and policy compliance classification
 521 (Kumar et al., 2024). These methods assess veracity against external knowledge, evaluate internal
 522 consistency, or identify hate speech, harassment, and other violations based on semantics rather than
 523 surface keywords. Complementary works study robustness through out-of-distribution detection, e.g.,
 524 by adjusting likelihoods from generative models with input complexity measures to better separate
 525 in- and out-of-distribution data (Serrà et al., 2020). While effective for detecting policy violations,
 526 factual errors, or distributional anomalies, these approaches typically focus on discrete factuality
 527 (“Is this claim true?”) or policy compliance (“Does this post violate rules?”) and do not address the
 528 distributional aspect of deception as we define it, at least not directly. Our work aims to fill this gap.

530 8 CONCLUSION

531 This work offers a new perspective on media integrity by shifting the focus from detecting individual
 532 deepfakes to identifying broader semantic distortions. As generative models improve, low-level
 533 artifacts may disappear, but the risk of misleading content remains. We propose semantic calibration
 534 as a lightweight, interpretable method that complements existing deepfake detectors by targeting
 535 distributional shifts explicitly. Beyond deception reduction, this approach could also be used for
 536 alternative goals such as personalized filtering, balanced media streams, scalable moderation, and
 537 more. It lays the groundwork for information systems that prioritize what content communicates,
 538 rather than how it was created.

540 REPRODUCIBILITY STATEMENT
541

542 All datasets used in our experiments (see Section 4) are publicly available and can be accessed through
543 our HuggingFace repository (Appendix A.7.2), with preprocessing details provided in Appendix A.5.
544 The prompts used for the captioning models are listed in Appendix A.6. All scripts, including those
545 for LLM fine-tuning, captioning, and rejection sampling, are available in our GitHub repository
546 (Appendix A.7.1), which also contains all hyperparameter configurations. Compute resources and
547 hardware details are reported in Appendix A.8. Finally, the fine-tuned models corresponding to each
548 experiment are released via our HuggingFace repository (Appendix A.7.2).

549
550 REFERENCES

551 S. Agarwal and L. R. Varshney. Limits of Deepfake Detection: A Robust Estimation Viewpoint. In *Proceedings*
552 of the ICML Workshop on Deep Learning for Detecting AudioVisual Fakes, 2019.

553 Hadeer Ahmed, Issa Traore, and Sherif Saad. Detection of online fake news using n-gram analysis and
554 machine learning techniques. In *Intelligent, Secure, and Dependable Systems in Distributed and Cloud*
555 *Environments (ISDDC 2017)*, volume 10618 of *Lecture Notes in Computer Science*, pp. 127–138. Springer,
556 Cham, 2017. doi: 10.1007/978-3-319-69155-8_9.

557 Fabian Caba Heilbron, Victor Escorcia, Bernard Ghanem, and Juan Carlos Niebles. Activitynet: A large-scale
558 video benchmark for human activity understanding. In *Proceedings of the IEEE conference on computer*
559 *vision and pattern recognition*, pp. 961–970, 2015.

560 Honglie Chen, Weidi Xie, Andrea Vedaldi, and Andrew Zisserman. VGGSound: A Large-Scale Audio-Visual
561 Dataset. In *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp.
562 721–725, 2020.

563 Kanzhi Cheng, Wenpo Song, Jiaxin Fan, Zheng Ma, Qiushi Sun, Fangzhi Xu, Chenyang Yan, Nuo Chen,
564 Jianbing Zhang, and Jiajun Chen. CapArena: Benchmarking and analyzing detailed image captioning in
565 the LLM era. In *Findings of the Association for Computational Linguistics: ACL 2025*, pp. 14077–14094.
566 Association for Computational Linguistics, 2025.

567 Yunfei Chu, Jin Xu, Xiaohuan Zhou, Qian Yang, Shiliang Zhang, Zhijie Yan, Chang Zhou, and Jingren Zhou.
568 Qwen-Audio: Advancing Universal Audio Understanding via Unified Large-Scale Audio-Language Models.
569 *arXiv preprint arXiv:2311.07919*, 2023.

570 Riccardo Corvi, Davide Cozzolino, Giada Zingarini, Giovanni Poggi, Koki Nagano, and Luisa Verdoliva. On
571 the Detection of Synthetic Images Generated by Diffusion Models. In *IEEE International Conference on*
572 *Acoustics, Speech and Signal Processing (ICASSP)*, pp. 1–5, 2023.

573 Thomas M. Cover and Joy A. Thomas. *Elements of Information Theory*. Wiley-Interscience, 2006.

574 Charles Day. The Future of Misinformation. *Computing in Science and Engineering*, 21(1):108, 2019.

575 Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Fei-Fei Li. ImageNet: A Large-Scale Hierarchical
576 Image Database. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*
577 (*CVPR*), pp. 248–255, 2009.

578 Brian Dolhansky, Joanna Bitton, Ben Pflaum, Jikuo Lu, Russ Howes, Menglin Wang, and Cristian Canton Ferrer.
579 The Deepfake Detection Challenge (DFDC) Dataset. *arXiv preprint arXiv:2006.07397*, 2020.

580 Ricard Durall, Margret Keuper, and Janis Keuper. Watch Your Up-Convolution: CNN Based Image Manipulation
581 Detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*
582 *Workshops (CVPRW)*, pp. 1900–1909, 2020.

583 Joel Frank, Thorsten Eisenhofer, Lea Schönherr, Asja Fischer, Dorothea Kolossa, and Thorsten Holz. Leveraging
584 Frequency Analysis for Deep Fake Image Recognition. In *Proceedings of the 37th International Conference*
585 *on Machine Learning (ICML)*, pp. 3247–3258, 2020.

586 Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron
587 Courville, and Yoshua Bengio. Generative Adversarial Nets. In *Proceedings of the 27th Conference on*
588 *Neural Information Processing Systems (NIPS)*, 2014.

589 Arash Heidari, Nima Jafari Navimipour, Hasan Dag, and Mehmet Unal. Deepfake Detection Using Deep
590 Learning Methods: A Systematic and Comprehensive Review. *Wiley Interdisciplinary Reviews: Data Mining*
591 and *Knowledge Discovery*, 14(2):e1520, 2024.

594 Ari Holtzman, Jan Buys, Li Du, Maxwell Forbes, and Yejin Choi. The Curious Case of Neural Text Degeneration.
 595 In *Proceedings of the 8th International Conference on Learning Representations (ICLR)*, 2020.

596

597 Beizhe Hu, Qiang Sheng, Juan Cao, Yuhui Shi, Yang Li, Danding Wang, and Peng Qi. Bad Actor, Good Advisor:
 598 Exploring the Role of Large Language Models in Fake News Detection. In *Proceedings of the 38th AAAI
 599 Conference on Artificial Intelligence (AAAI)*, pp. 22105–22113, 2024.

600

601 Shehzeen Hussain, Paarth Neekhara, Malhar Jere, Farinaz Koushanfar, and Julian McAuley. Adversarial
 602 Deepfakes: Evaluating Vulnerability of Deepfake Detectors to Adversarial Examples. In *Proceedings of the
 603 IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)*, pp. 3468–3477, 2021.

604

605 Muhammad U. Ikram, Sohail J. Abuadbba, M. Ali Akber, Athar Ali, Kristen Moore, Josef Pieprzyk, and Simon S.
 606 Woo. Can Deepfake Detectors Be Trusted? A Benchmark and Robustness Analysis. *IEEE Transactions on
 607 Information Forensics and Security*, 2024.

608

609 Douwe Kiela, Hamed Firooz, Aravind Mohan, Vedanuj Goswami, Amanpreet Singh, Pratik Ringshia, and
 610 Davide Testuggine. The hateful memes challenge: Detecting hate speech in multimodal memes, 2021.

611

612 Neema Kotonya and Francesca Toni. Explainable Automated Fact-Checking: A Survey. In *Proceedings of the
 613 28th International Conference on Computational Linguistics (COLING)*, pp. 5430–5443, 2020.

614

615 Ranjay Krishna, Kenji Hata, Frederic Ren, Li Fei-Fei, and Juan Carlos Niebles. Dense-captioning events in
 616 videos. In *Proceedings of the International Conference on Computer Vision (ICCV)*, 2017.

617

618 KristinCoyote. Us senators twitter data. [https://www.kaggle.com/datasets/kristincoyote/
 619 us-senators-twitter-data](https://www.kaggle.com/datasets/kristincoyote/us-senators-twitter-data), 2025. [Dataset, accessed Nov 20].

620

621 Alex Krizhevsky. Learning Multiple Layers of Features from Tiny Images. Technical report, University of
 622 Toronto, 2009. CIFAR-10 and CIFAR-100 datasets.

623

624 Deepak Kumar, Yousef Anees AbuHashem, and Zakir Durumeric. Watch Your Language: Investigating Content
 625 Moderation with Large Language Models. In *Proceedings of the 18th International AAAI Conference on Web
 626 and Social Media (ICWSM)*, pp. 865–878, 2024.

627

628 Christian Ledig, Lucas Theis, Ferenc Huszár, Jose Caballero, Andrew Cunningham, Alejandro Acosta, Andrew
 629 Aitken, Alykhan Tejani, Johannes Totz, Zehan Wang, and Wenzhe Shi. Photo-Realistic Single Image Super-
 630 Resolution Using a Generative Adversarial Network. In *Proceedings of the IEEE Conference on Computer
 631 Vision and Pattern Recognition (CVPR)*, pp. 4681–4690, 2017.

632

633 Alyssa Lees, Siva Reddy Sunkara, Kiat Wee Zhang, Ala Bsharat, Marcos Zampieri, Paula Fortuna, John
 634 Pavlopoulos, Jeffrey Sorensen, Lucas Dixon, and Nithum Thain. A unified multilingual approach to toxic
 635 content classification. In *Proceedings of the 60th Annual Meeting of the Association for Computational
 636 Linguistics (ACL)*. Association for Computational Linguistics, 2022.

637

638 Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and
 639 C. Lawrence Zitnick. Microsoft COCO: Common Objects in Context. In *European Conference on Computer
 640 Vision (ECCV)*, pp. 740–755, 2014a.

641

642 Tsung-Yi Lin, Michael Maire, Serge J. Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and
 643 C. Lawrence Zitnick. Microsoft COCO: Common Objects in Context. In *Proceedings of the 13th European
 644 Conference on Computer Vision (ECCV)*, pp. 740–755, 2014b.

645

646 Ilya Loshchilov and Frank Hutter. Decoupled Weight Decay Regularization. In *Proceedings of the 7th
 647 International Conference on Learning Representations (ICLR)*, 2019.

648

649 Yuan Mei, Mengchen Liu, Hakan Bilen, Di Tu, Ming Bai, Weidi Xie, and Andrew Zisserman. AudioSetCaps:
 650 Generating Captions for AudioSet. In *IEEE Winter Conference on Applications of Computer Vision (WACV)*,
 651 pp. 6243–6252, 2023.

652

653 Fabian Mentzer, George D. Toderici, Michael Tschanne, and Eirikur Agustsson. High-Fidelity Generative
 654 Image Compression. In *Proceedings of the 33rd Conference on Neural Information Processing Systems
 655 (NeurIPS)*, pp. 11913–11924, 2020.

656

657 Yisroel Mirsky and Wenke Lee. The Creation and Detection of Deepfakes: A Survey. *ACM Computing Surveys*,
 658 54(1):1–41, 2021.

659

660 nbandhi. Political podcasts listing with audio links. [https://www.kaggle.com/datasets/nbandhi/
 661 political-podcasts-listing-with-audio-links](https://www.kaggle.com/datasets/nbandhi/political-podcasts-listing-with-audio-links), 2025. [Dataset, accessed Nov 20].

648 Britt Paris and Joan Donovan. Deepfakes and Cheap Fakes: The Manipulation of Audio and Visual Evidence.
 649 Technical report, Data & Society Research Institute, 2019.
 650

651 Project Gutenberg. Project Gutenberg, 2025. URL <https://www.gutenberg.org>. An online library of
 652 free eBooks. Accessed 2025.
 653

654 Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. Language Models are
 655 Unsupervised Multitask Learners, 2019. OpenAI Blog, GPT-2 model release.
 656

657 Aditya Ramesh, Prafulla Dhariwal, Alex Nichol, Casey Chu, and Mark Chen. Hierarchical Text-Conditional
 658 Image Generation with CLIP Latents, 2022. arXiv preprint arXiv:2204.06125.
 659

660 Md Shohel Rana, Mohammad Nur Nobi, Beddhu Murali, and Andrew H. Sung. Deepfake Detection: A
 661 Systematic Literature Review. *IEEE Access*, 10:25494–25513, 2022.
 662

663 Reve.art. Reve.art, 2025. URL <https://reve.art>. Accessed: April 15th 2025.
 664

665 Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-Resolution
 666 Image Synthesis with Latent Diffusion Models. In *Proceedings of the IEEE/CVF Conference on Computer
 667 Vision and Pattern Recognition (CVPR)*, pp. 10684–10695, 2022.
 668

669 Andreas Rössler, Davide Cozzolino, Luisa Verdoliva, Christian Riess, Justus Thies, and Matthias Nießner. Face-
 670 Forensics++: Learning to Detect Manipulated Facial Images. In *Proceedings of the IEEE/CVF International
 671 Conference on Computer Vision (ICCV)*, pp. 1–11, 2019.
 672

673 Chitwan Saharia, William Chan, Saurabh Saxena, Lala Li, Jay Whang, Emily L. Denton, Kamyar Ghasemipour,
 674 Burcu Karagol Ayan, S. Sara Mahdavi, Rapha Gontijo Lopes, Tim Salimans, Jonathan Ho, David J. Fleet, and
 675 Mohammad Norouzi. Photorealistic Text-to-Image Diffusion Models with Deep Language Understanding. In
 676 *Proceedings of the 35th Conference on Neural Information Processing Systems (NeurIPS)*, pp. 36479–36494,
 677 2022.
 678

679 Chitwan Saharia, Jonathan Ho, William Chan, Tim Salimans, David J. Fleet, and Mohammad Norouzi. Image
 680 Super-Resolution via Iterative Refinement. *IEEE Transactions on Pattern Analysis and Machine Intelligence*,
 681 45(7):9088–9100, 2023.
 682

683 Justin Salamon, Christopher Jacoby, and Juan Pablo Bello. A Dataset and Taxonomy for Urban Sound Research.
 684 In *ACM International Conference on Multimedia (ACM MM)*, pp. 1041–1044, 2014.
 685

686 Joan Serrà, Santiago Pascual, and Jordi Valle. Input complexity and out-of-distribution detection with likelihood-
 687 based generative models. In *International Conference on Learning Representations (ICLR)*, 2020.
 688

689 Claude E. Shannon. Prediction and Entropy of Printed English. *The Bell System Technical Journal*, 30(1):50–64,
 690 1951.
 691

692 Gowthami Somepalli, Vasu Singla, Micah Goldblum, Jonas Geiping, and Tom Goldstein. Diffusion Art or
 693 Digital Forgery? Investigating Data Replication in Diffusion Models. In *Proceedings of the IEEE/CVF
 694 Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 6048–6058, 2023a.
 695

696 Gowthami Somepalli, Vasu Singla, Micah Goldblum, Jonas Geiping, and Tom Goldstein. Understanding and
 697 Mitigating Copying in Diffusion Models. In *Proceedings of the 36th Conference on Neural Information
 698 Processing Systems (NeurIPS)*, pp. 47783–47803, 2023b.
 699

700 surajkarakulath. Labelled corpus political bias (hugging face). [https://www.kaggle.com/datasets/
 701 surajkarakulath/labelled-corpus-political-bias-hugging-face](https://www.kaggle.com/datasets/surajkarakulath/labelled-corpus-political-bias-hugging-face), 2025. [Dataset, accessed Nov 20].
 702

703 Luisa Verdoliva. Media Forensics and Deepfakes: An Overview. *IEEE Journal of Selected Topics in Signal
 704 Processing*, 14(5):910–932, 2020.
 705

706 Ivan Vykopal, Matúš Pikuliak, Simon Ostermann, and Marián Šimko. Generative Large Language Models in
 707 Automated Fact-Checking: A Survey. *arXiv preprint arXiv:2407.02351*, 2024.
 708

709 Sheng-Yu Wang, Oliver Wang, Richard Zhang, Andrew Owens, and Alexei A. Efros. CNN-Generated Images
 710 Are Surprisingly Easy to Spot... For Now. In *Proceedings of the IEEE/CVF Conference on Computer Vision
 711 and Pattern Recognition (CVPR)*, pp. 8695–8704, 2020.
 712

713 William Yang Wang. “liar, liar pants on fire”: A new benchmark dataset for fake news detection. In *Proceedings
 714 of the 55th Annual Meeting of the Association for Computational Linguistics (Short Papers)*, ACL 2017, pp.
 715 422–426, Vancouver, Canada, 2017. Association for Computational Linguistics. doi: 10.18653/v1/P17-2067.
 716

702 Bangzhen Wen, Yiling Xu, Hongxia Wang, Han Su, Haoliang Li, Alex C. Kot, and Kwok-Yan Lam. Deepfake
703 Detection: A Reliability-Aware Survey. *arXiv preprint arXiv:2211.14001*, 2022.

704

705 Junxian Wu, Yuting Huang, Zeyu Li, Hanlin Wang, Yuxiang Zhang, Yichong Zhou, Kun Xu, and Yiran Chen.
706 Icm-assistant: Instruction-tuned multimodal llms for intelligent content moderation with explanation. *arXiv
707 preprint arXiv:2412.18216*, 2024.

708

709 Jiahui Yuan, Zhiheng Xu, Yikai Wang, Xinyuan Chen, Yuejie Tang, Jie Zhang, and Junchi Xu. Asymmetric
710 mixed-modal moderation (am3): Cross-modal learning for multimodal content moderation. In *Proceedings of
711 the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 19573–19582. IEEE,
712 2023.

713

714 Lu Yuan, Dong Chen, Yi-Ling Chen, Noel Codella, Xiyang Dai, Jianfeng Gao, Houdong Hu, Xuedong Huang,
715 Boxin Li, Chunyuan Li, Ce Liu, Mengchen Liu, Zicheng Liu, Pengchuan Luo, Dian Mei, Shuai Shi, Jianfeng
716 Sun, Jian Wang, Yuhang Wang, Zhe Wang, Wenhui Wang, Bin Xiao, Lu Wang, Furu Wei, Ting Yao, Jianqiang
717 Yu, Zhenting Yue, Jianwei Zhang, Zhe Zhang, Yang Zhang, Yi Zhang, Dejing Zhou, and Lijuan Zhu. Florence:
718 A New Foundation Model for Computer Vision, 2021. *arXiv preprint arXiv:2111.11432*.

719

720 Xiang Zhang, Junbo Zhao, and Yann LeCun. Character-Level Convolutional Networks for Text Classification.
721 In *Proceedings of the 28th Conference on Neural Information Processing Systems (NIPS)*, pp. 649–657, 2015.

722

723 Gonzalo A. Álvarez Hervás. Fake news tfg dataset. Dataset at Hugging Face, 2022.

724

725

726

727

728

729

730

731

732

733

734

735

736

737

738

739

740

741

742

743

744

745

746

747

748

749

750

751

752

753

754

755

756

A APPENDIX

757

A.1 SEMANTIC VS NON-SEMANTIC INFORMATION

760 We estimate how much of an image’s information
 761 is semantic. Consider a small 256×256
 762 RGB image compressed near the perceptual
 763 threshold (≈ 0.3 bits per pixel, as achieved by
 764 state-of-the-art methods Mentzer et al. (2020)),
 765 and suppose the semantic content of such an
 766 image can be expressed in 100 English words
 767 (this is a conservative estimate, given that the
 768 average caption in human-annotated datasets like
 769 MS-COCO Lin et al. (2014b) is only about 10
 770 words). Using Shannon’s estimate of roughly
 771 1.3 bits per character and an average of 4.5 char-
 772 acters per word Shannon (1951), this amounts to
 773 ≈ 600 bits of semantic information, only about
 $600/(256^2 \times 0.3) \approx 3\%$ of the image’s total perceptual information. Figure 6 shows two synthetic images
 774 that convey the same semantic information:

775 *A male lion lying on open savannah grassland, resting under a tree positioned to its right. Additional
 776 trees are visible in the distant background. (24 words)*

777 While minor differences are visible (e.g., the lion’s head orientation or the exact distance of back-
 778 ground trees), these can also be considered non-semantic, as they would likely go unnoticed if only
 779 one image were shown. Indeed, a human asked to caption a single image in Figure 6 would typically
 780 not mention such details, as they are not central to the perceived meaning of the scene. Similarly, one
 781 would most likely not specify the orientation of the lion’s head while prompting a generative model.

782

A.2 NOTATION

783

A.2.1 VARIABLES

784 We define the following random variables shown in Figure 2:

- 785 • $X \in \mathcal{X}$: Observed variable representing the media of interest (e.g., text, images, video,
 786 audio, etc.).
- 787 • $Z \in \mathcal{Z}$: Latent variable representing all the semantic information present in a media X .
- 788 • $\hat{Z} \in \hat{\mathcal{Z}}$: Observed variable representing the textual representation of the semantic informa-
 789 tion present in a media X , extracted using a deterministic captioning model \hat{f} such that
 $\hat{Z} = \hat{f}(X)$.
- 790 • $G \in \{0, 1\}$: Latent variable indicating if X is generated ($G = 1$) or is real ($G = 0$).
- 791 • $F_\phi \in \{0, 1\}$: Observed variable indicating whether X is flagged ($F_\phi = 1$) or not ($F_\phi = 0$).
 792 The subscript ϕ indicates the dependence on the decision rule $\phi(x) \triangleq \mathbb{P}[F_\phi = 1 | X = x]$.
- 793 • $E_\phi \in \{0, 1\}$: Latent error variable (i.e., $E = \mathbb{1}\{F_\phi \neq G\}$) with $\mathbb{1}$ the indicator variable.

801

A.2.2 DISTRIBUTIONS

802 Let

$$803 p(g, z, x, \hat{g}, e) = \mathbb{P}[G = g, Z = z, X = x, \hat{G} = \hat{g}, E = e]$$

804 denote the joint probability distribution of the above random variables. When the context is clear,
 805 we will also use p to refer to the marginal distribution over any subset of these variables (e.g.,
 $p(x) = \mathbb{P}[X = x]$). We denote conditioning on $G = 1$ (generated media) by using q in place of p ,
 806 and conditioning on $G = 0$ (real media) by using p_r instead of p . In addition, we use the superscript
 p^ϕ to indicate conditioning on $F_\phi = 0$ (filtered stream). More precisely:

- $p_r(x) = \mathbb{P}[X = x | G = 0]$: The distribution of the real media samples.
- $p_r(x|z) = \mathbb{P}[X = x | Z = z, G = 0]$: The distribution of real media that have semantic information z .
- $q(x) = \mathbb{P}[X = x | G = 0]$: The distribution of the fake media samples.
- $q(x|z) = \mathbb{P}[X = x | Z = z, G = 1]$: The distribution induced by the conditional generative model, typically a neural network (we assume that we can sample from it), for media samples that have semantic information z .
- $p_r(z) = \mathbb{P}[Z = z | G = 0]$: The distribution of latent semantics information found in the real media samples.
- $q(z) = \mathbb{P}[Z = z | G = 1]$: The distribution of latent semantic information specified by users when generating a media.
- $p^\phi(z) = \mathbb{P}[Z = z | F_\phi = 0]$: The distribution of latent semantics information found in the filtered media samples.
- $p_r^\phi(z) = \mathbb{P}[Z = z | G = 0, F_\phi = 0]$: The distribution of latent semantics information of the real media samples found in the filtered media stream.
- $q^\phi(z) = \mathbb{P}[Z = z | G = 1, F_\phi = 0]$: The distribution of latent semantics information of the generated media samples found in the filtered media stream.

In addition, we define $p_G \triangleq \mathbb{P}[G = 1]$ and $p_G^\phi \triangleq \mathbb{P}[G = 1 | F = 0]$ to shorten notation, and we recall the definition of the conditional entropy

$$H(Z | \hat{Z}) = - \sum_{\hat{z} \in \hat{\mathcal{Z}}} \mathbb{P}[\hat{Z} = \hat{z}] \sum_{z \in \mathcal{Z}} \mathbb{P}[Z = z | \hat{Z} = \hat{z}] \log \mathbb{P}[Z = z | \hat{Z} = \hat{z}].$$

A.3 FORWARD OBJECTIVE

The objective in Eq. equation 2 is defined using the *backward* KL divergence. An alternative (but arguably harder) training objective is to minimize the *forward* KL divergence $\mathcal{L}(p_r, q)$, which encourages the generator to replicate *all* the variability present in real media. However, for our analysis, Theorem 1 also holds if one assume $\mathcal{L}(p_r, q) \leq \epsilon$ instead of $\mathcal{L}(q, p_r) \leq \epsilon$. For clarity, we only focus on the reverse KL, as it does not penalize the generator for missing low-probability modes of p_r , and thus tends to be easier to minimize. Indeed, training a generator that only generates convincing media samples is arguably easier than training a generator that can generate *any* type of real media.

A.4 PROOFS

This section presents the derivations supporting the main results of the paper.

A.4.1 CONVEXITY OF KL DIVERGENCE

We first recall that KL divergence is convex in its first argument for discrete distributions Cover & Thomas (2006). Let p_1 , p_2 , and p_3 be distributions over a discrete space \mathcal{Z} , and let $\lambda \in [0, 1]$ such that $p_\lambda(z) = \lambda p_1(z) + (1 - \lambda)p_2(z)$. We have

$$\begin{aligned} D_{\text{KL}}(p_\lambda(z) \| p_3(z)) &= \sum_{z \in \mathcal{Z}} p_\lambda(z) \log \frac{p_\lambda(z)}{p_3(z)} \\ &= \sum_{z \in \mathcal{Z}} [\lambda p_1(z) + (1 - \lambda)p_2(z)] \log \frac{\lambda p_1(z) + (1 - \lambda)p_2(z)}{p_3(z)} \\ &\leq \lambda \sum_{z \in \mathcal{Z}} p_1(z) \log \frac{p_1(z)}{p_3(z)} + (1 - \lambda) \sum_{z \in \mathcal{Z}} p_2(z) \log \frac{p_2(z)}{p_3(z)} \\ &= \lambda D_{\text{KL}}(p_1(z) \| p_3(z)) + (1 - \lambda) D_{\text{KL}}(p_2(z) \| p_3(z)), \end{aligned}$$

864 where the inequality follows from Jensen’s inequality applied pointwise to the convex function
 865 $x \mapsto x \log(x/p_3(z))$. We can now prove the desired inequality. From the definition of p^ϕ in
 866 Eq. equation 4, we have

$$867 \quad p^\phi(z) = p_G^\phi q^\phi(z) + (1 - p_G^\phi) p_r^\phi(z).$$

869 Applying the convexity of KL divergence with $\lambda = p_G^\phi$, $p_1 = q^\phi$ and $p_2 = p_r^\phi$, we obtain

$$871 \quad D_{\text{KL}}(p^\phi(z) \parallel p_r(z)) \leq p_G^\phi D_{\text{KL}}(q^\phi(z) \parallel p_r(z)) + (1 - p_G^\phi) D_{\text{KL}}(p_r^\phi(z) \parallel p_r(z)).$$

872 By definition of δ in Eq. equation 3, this yields

$$874 \quad \delta(p^\phi \mid p_r) \leq p_G^\phi \delta(q^\phi \mid p_r) + (1 - p_G^\phi) \delta(p_r^\phi \mid p_r).$$

876 A.4.2 DIVERGENCE CHAIN RULE

878 We now establish the following intermediate result that will be instrumental in proving Theorem 1:

$$879 \quad D_{\text{KL}}(q(x) \parallel p_r(x)) \leq \mathcal{L}(q, p_r) + D_{\text{KL}}(q(z) \parallel p_r(z)). \quad (10)$$

881 Starting from the definition, and observing that $q(x) = q(z, x)$ and $p_r(x) = p_r(z, x)$ for $z = f(x)$
 882 (since f is deterministic, meaning each media item x maps to a unique semantic representation
 883 $z = f(x)$), and letting $\mathcal{X}(z) \triangleq x \in \mathcal{X} \mid f(x) = z$, we have

$$\begin{aligned} 885 \quad D_{\text{KL}}(q(x) \parallel p_r(x)) &= \sum_{x \in \mathcal{X}} q(x) \log \frac{q(x)}{p_r(x)} \\ 886 \\ 887 &= \sum_{z \in \mathcal{Z}} \sum_{x \in \mathcal{X}(z)} q(x) \log \frac{q(z, x)}{p_r(z, x)} \\ 888 \\ 889 &= \sum_{z \in \mathcal{Z}} \sum_{x \in \mathcal{X}(z)} q(z) q(x|z) \log \frac{q(z) q(x|z)}{p_r(z) p_r(x|z)} \\ 890 \\ 891 &= \sum_{z \in \mathcal{Z}} \sum_{x \in \mathcal{X}(z)} q(z) q(x|z) \log \frac{q(z)}{p_r(z)} + \sum_{z \in \mathcal{Z}} \sum_{x \in \mathcal{X}(z)} q(z) q(x|z) \log \frac{q(x|z)}{p_r(x|z)} \\ 892 \\ 893 &= \sum_{z \in \mathcal{Z}} q(z) \log \frac{q(z)}{p_r(z)} \left(\sum_{x \in \mathcal{X}(z)} q(x|z) \right) + \sum_{z \in \mathcal{Z}} q(z) \left(\sum_{x \in \mathcal{X}(z)} q(x|z) \log \frac{q(x|z)}{p_r(x|z)} \right) \\ 894 \\ 895 &= \sum_{z \in \mathcal{Z}} q(z) \log \frac{q(z)}{p_r(z)} \left(\sum_{x \in \mathcal{X}} q(x|z) \right) + \sum_{z \in \mathcal{Z}} q(z) \left(\sum_{x \in \mathcal{X}} q(x|z) \log \frac{q(x|z)}{p_r(x|z)} \right) \\ 896 \\ 897 &= \sum_{z \in \mathcal{Z}} q(z) \log \frac{q(z)}{p_r(z)} + \sum_{z \in \mathcal{Z}} q(z) D_{\text{KL}}(q(x|z) \parallel p_r(x|z)) \\ 898 \\ 899 &= D_{\text{KL}}(q(z) \parallel p_r(z)) + \mathbb{E}_{z \sim q(z)} [D_{\text{KL}}(q(x|z) \parallel p_r(x|z))], \end{aligned}$$

900 where we used the fact that $q(x \mid z) = p_r(x \mid z) = 0$ for $z \neq f(x)$, and adopted the standard
 901 information-theoretic convention $0 \log(\frac{0}{0}) \equiv 0$ to extend the sums from $\mathcal{X}(z)$ to \mathcal{X} . By definition,
 902 we have

$$910 \quad \mathbb{E}_{z \sim q(z)} [D_{\text{KL}}(q(x|z) \parallel p_r(x|z))] \leq \sup_{z \in \mathcal{Z}} D_{\text{KL}}(q(x|z) \parallel p_r(x|z)) = \mathcal{L}(q, p_r).$$

912 Combining both results, we directly obtain Eq. 10.

914 A.4.3 LIMIT OF DEEPFAKE DETECTION

916 First, we can express the probability of error of any decision rule ϕ as

$$917 \quad \mathbb{P}[E_\phi = 1] = \mathbb{P}[F_\phi = 0, G = 1] + \mathbb{P}[F_\phi = 1, G = 0]. \quad (11)$$

918 Using the factorization provided in Figure 2, we can expand each term as follows:
919

$$\begin{aligned}
920 \quad \mathbb{P}[F_\phi = 0, G = 1] &= \sum_{x \in \mathcal{X}} \mathbb{P}[F_\phi = 0, G = 1, X = x] \\
921 \\
922 &= \sum_{x \in \mathcal{X}} \mathbb{P}[G = 1] \mathbb{P}[X = x | G = 1] \mathbb{P}[F_\phi = 0, X = x] \\
923 \\
924 &= \sum_{x \in \mathcal{X}} p_G q(x) (1 - \phi(x)). \\
925 \\
926 \\
927
\end{aligned}$$

928 Additionally,
929

$$\begin{aligned}
930 \quad \mathbb{P}[F_\phi = 1, G = 0] &= \sum_{x \in \mathcal{X}} \mathbb{P}[F_\phi = 1, G = 0, X = x] \\
931 \\
932 &= \sum_{x \in \mathcal{X}} \mathbb{P}[G = 0] \mathbb{P}[X = x | G = 0] \mathbb{P}[F_\phi = 1, X = x] \\
933 \\
934 &= \sum_{x \in \mathcal{X}} (1 - p_G) p_r(x) \phi(x). \\
935 \\
936 \\
937
\end{aligned}$$

938 Substituting both expressions in Eq. equation 11:
939

$$\begin{aligned}
940 \quad \mathbb{P}[E_\phi = 1] &= \mathbb{P}[F_\phi = 0, G = 1] + \mathbb{P}[F_\phi = 1, G = 0] \\
941 \\
942 &= \sum_{x \in \mathcal{X}} [p_G q(x) (1 - \phi(x)) + (1 - p_G) p_r(x) \phi(x)] \\
943 \\
944 &\geq \sum_{x \in \mathcal{X}} \min\{p_G q(x), (1 - p_G) p_r(x)\}. \\
945 \\
946
\end{aligned} \tag{12}$$

947 Recall that $a + b - |b - a| = 2 \min\{a, b\}$ for any $a, b \geq 0$. Given any x, z and setting
948

$$a = p_G q(x) \quad \text{and} \quad b = (1 - p_G) p_r(x),$$

949 we can write
950

$$951 \quad \min\{p_G q(x), (1 - p_G) p_r(x)\} = \frac{1}{2} [p_G q(x) + (1 - p_G) p_r(x) - |(1 - p_G) p_r(x) - p_G q(x)|].$$

952 Summing over x , we obtain
953

$$954 \quad \sum_{x \in \mathcal{X}} \min\{p_G q(x), (1 - p_G) p_r(x)\} = \frac{1}{2} \left[1 - \sum_{x \in \mathcal{X}} |(1 - p_G) p_r(x) - p_G q(x)| \right]. \tag{13}$$

955 Next, note that the total variation distance between two distributions $p(x)$ and $q(x)$ is defined as
956 $\text{TV}[p(x), q(x)] = \frac{1}{2} \sum_{x \in \mathcal{X}} |p(x) - q(x)|$. The key step is to relate the sum in the right-hand side of
957 Eq. equation 13 to $\text{TV}[p(x), q(x)]$. To that end, we identify two cases:
958

- 959 • If $p_G \leq 1 - p_G$ (i.e., $p_G \leq \frac{1}{2}$), then

$$\begin{aligned}
960 \quad |(1 - p_G) p_r(x) - p_G q(x)| &= |p_G(p_r(x) - q(x)) + (1 - 2p_G) p_r(x)| \\
961 \\
962 &\leq |p_G(p_r(x) - q(x))| + |(1 - 2p_G) p_r(x)| \\
963 \\
964 &= p_G |p_r(x) - q(x)| + (1 - 2p_G) p_r(x),
\end{aligned}$$

965 where we have used the triangular inequality. Summing over x :
966

$$\begin{aligned}
967 \quad \sum_{x \in \mathcal{X}} |(1 - p_G) p_r(x) - p_G q(x)| &\leq \sum_{x \in \mathcal{X}} p_G |p_r(x) - q(x)| + \sum_{x \in \mathcal{X}} (1 - 2p_G) p_r(x) \\
968 \\
969 &= 2p_G \text{TV}[p_r(x), q(x)] + (1 - 2p_G).
\end{aligned}$$

972
973

Therefore:

974
975
976
977
978
979
980
981

$$\begin{aligned} \sum_{x \in \mathcal{X}} \min\{p_G q(x), (1 - p_G) p_r(x)\} &= \frac{1}{2} \left[1 - \sum_{x \in \mathcal{X}} |(1 - p_G) p_r(x) - p_G q(x)| \right] \\ &\geq \frac{1}{2} [1 - 2p_G \text{TV}[p_r(x), q(x)] - (1 - 2p_G)] \\ &= p_G - p_G \text{TV}[p_r(x), q(x)] \\ &= \min\{p_G, 1 - p_G\} (1 - \text{TV}[p_r(x), q(x)]). \end{aligned}$$

982
983

- If $p_G \geq 1 - p_G$ (i.e., $p_G \geq \frac{1}{2}$), then

984
985
986
987

$$\begin{aligned} |(1 - p_G) p_r(x) - p_G q(x)| &= |(1 - p_G) [p_r(x) - q(x)] + (1 - 2p_G) q(x)| \\ &\leq |(1 - p_G) [p_r(x) - q(x)]| + |(1 - 2p_G) q(x)| \\ &= (1 - p_G) |p_r(x) - q(x)| + (2p_G - 1) q(x). \end{aligned}$$

988
989Summing over x :990
991
992
993
994
995
996

$$\begin{aligned} \sum_{x \in \mathcal{X}} |(1 - p_G) p_r(x) - p_G q(x)| &= \sum_{x \in \mathcal{X}} |(1 - p_G) p_r(x) - p_G q(x)| \\ &\leq \sum_{x \in \mathcal{X}} (1 - p_G) |p_r(x) - q(x)| + \sum_{x \in \mathcal{X}} (2p_G - 1) q(x) \\ &= 2(1 - p_G) \text{TV}[p_r(x), q(x)] + (2p_G - 1). \end{aligned}$$

997
998

Therefore:

999
1000
1001
1002
1003
1004
1005
1006

$$\begin{aligned} \sum_{x \in \mathcal{X}} \min\{p_G q(x), (1 - p_G) p_r(x)\} &= \frac{1}{2} \left[1 - \sum_{x \in \mathcal{X}} |(1 - p_G) p_r(x) - p_G q(x)| \right] \\ &\geq \frac{1}{2} [1 - 2(1 - p_G) \text{TV}[p_r(x), q(x)] - (2p_G - 1)] \\ &= (1 - p_G) - (1 - p_G) \text{TV}[p_r(x), q(x)] \\ &= \min\{p_G, 1 - p_G\} (1 - \text{TV}[p_r(x), q(x)]). \end{aligned}$$

1007
1008

Since both cases yield the same bound, we can substitute it in Eq. equation 12 to obtain

1009
1010

$$\inf_{\phi} \mathbb{P}[E_{\phi} = 1] \geq \min\{p_G, 1 - p_G\} (1 - \text{TV}[p_r(x), q(x)]).$$

1011
1012
1013Using Pinsker's inequality $\text{TV}[p(x), q(x)] \leq \sqrt{\frac{1}{2} D_{\text{KL}}(q(x) \| p_r(x))}$, we can write1014
1015
1016
1017
1018
1019
1020
1021
1022
1023
1024

$$\begin{aligned} \inf_{\phi} \mathbb{P}[E_{\phi} = 1] &\geq \min\{p_G, 1 - p_G\} (1 - \text{TV}[p_r(x), q(x)]) \\ &\geq \min\{p_G, 1 - p_G\} \left(1 - \sqrt{\frac{1}{2} D_{\text{KL}}(q(x) \| p_r(x))} \right) \\ &\geq \min\{p_G, 1 - p_G\} \left(1 - \sqrt{\frac{\mathcal{L}(q, p_r) + D_{\text{KL}}(q(z) \| p_r(z))}{2}} \right), \\ &\geq \min\{p_G, 1 - p_G\} \left(1 - \sqrt{\frac{\epsilon + \delta(q \mid p_r)}{2}} \right), \end{aligned} \tag{14}$$

1025

where we substituted our result from Eq. equation 10 in the second-to-last step, and then used the assumption that $\mathcal{L}(q, p_r) \leq \epsilon$ along with the definition from Eq. equation 3. Finally, we conclude the

1026 proof as follows:
 1027

$$\begin{aligned}
 1028 \sup_{\phi} \mathbb{P}[E_{\phi} = 0] &= 1 - \inf_{\phi} \mathbb{P}[E_{\phi} = 1] \\
 1029 &\leq 1 - \min\{p_G, 1 - p_G\} \left(1 - \sqrt{\frac{\epsilon + \delta(q \mid p_r)}{2}}\right) \\
 1030 &= 1 - \min\{p_G, 1 - p_G\} + \min\{p_G, 1 - p_G\} \sqrt{\frac{\epsilon + \delta(q \mid p_r)}{2}} \\
 1031 &\leq \max\{p_G, 1 - p_G\} + \max\{p_G, 1 - p_G\} \sqrt{\frac{\epsilon + \delta(q \mid p_r)}{2}}.
 \end{aligned}$$

1032
 1033 Note that we could have derived a similar bound with the alternative objective $\mathcal{L}'(q, p_r)$ discussed
 1034 in Appendix A.3 (defined with the forward KL divergence) if we had used the alternative form of
 1035 Pinsker's inequality $\text{TV}[p(x), q(x)] \leq \sqrt{1/2 D_{\text{KL}}(p(x) \parallel q(x))}$ in Eq. equation 14.
 1036

1041 A.4.4 REJECTION SAMPLING

1042 We show that the semantics distribution of accepted samples $p^{\phi}(z)$ is exactly $p_r(z)$. Let $F \in \{0, 1\}$
 1043 be a random variable indicating if a sample with semantical information Z is flagged (i.e., rejected).
 1044 We have:
 1045

$$\begin{aligned}
 1046 p^{\phi}(z) &= \mathbb{P}[Z = z \mid F = 0] \\
 1047 &= \frac{\mathbb{P}[Z = z, F = 0]}{\mathbb{P}[F = 0]} \\
 1048 &= \frac{\mathbb{P}[Z = z] \mathbb{P}[F = 0 \mid Z = z]}{\sum_{z' \in \mathcal{Z}} \mathbb{P}[F = 0, Z = z']} \\
 1049 &= \frac{p(z) \cdot \frac{p_r(z)}{Mp(z)}}{\sum_{z' \in \mathcal{Z}} p(z') \cdot \frac{p_r(z')}{Mp(z')}} \\
 1050 &= \frac{p_r(z)}{\sum_{z' \in \mathcal{Z}} p_r(z')} \\
 1051 &= p_r(z)
 \end{aligned}$$

1052 Thus, rejection sampling produces samples exactly from the desired distribution $p_r(z)$. Note that
 1053 $\mathbb{P}[F = 0] = \frac{1}{M}$, meaning that, on average, only one out of every M media sample is accepted.
 1054

1055 A.4.5 UPPER BOUND FOR M^*

1056 Assuming that $p(z) > 0$ for any z and $p_G > 0$, and defining $\mathcal{Z}_+ \triangleq \{z \in \mathcal{Z} \mid p_r(z) > 0\}$, we can
 1057 derive an upper bound on M^* as follows:
 1058

$$\begin{aligned}
 1059 M^* &= \sup_{z \in \mathcal{Z}} \frac{p_r(z)}{p(z)} \\
 1060 &= \sup_{z \in \mathcal{Z}_+} \frac{p_r(z)}{p(z)} \\
 1061 &= \sup_{z \in \mathcal{Z}_+} \frac{p_r(z)}{p_G q(z) + (1 - p_G) p_r(z)} \\
 1062 &= \sup_{z \in \mathcal{Z}_+} \left(\frac{1}{(1 - p_G) + p_G \cdot \frac{q(z)}{p_r(z)}} \right) \\
 1063 &\leq \frac{1}{1 - p_G}.
 \end{aligned}$$

1064 This concludes the proof.
 1065

1080

A.4.6 RATIO APPROXIMATION

1081

1082 First, since we assume that the captioning model \hat{f} is such that $H(Z \mid \hat{f}(X)) = 0$, we know that there
 1083 exists a deterministic function $g : \hat{\mathcal{Z}} \rightarrow \mathcal{Z}$ such that $Z = g(\hat{f}(X))$. Assuming that the conditional
 1084 generative model perfectly captures reality (i.e., $q(x' \mid z) = p_r(x' \mid z)$ for any x', z), we have:

$$\begin{aligned}
 \frac{p_r(\hat{z})}{p(\hat{z})} &= \frac{\sum_{x' \in \mathcal{X}(\hat{z})} p_r(x')}{\sum_{x' \in \mathcal{X}(\hat{z})} p(x')} & \mathcal{X}(\hat{z}) &\triangleq \{x \mid \hat{f}(x) = \hat{z}\} \\
 &= \frac{\sum_{x' \in \mathcal{X}(\hat{z})} p_r(f(x'), x')}{\sum_{x' \in \mathcal{X}(\hat{z})} p(f(x'), x')} & p_{(r)}(x') &= p_{(r)}(z', x') \text{ for } z' \triangleq f(x') \\
 &= \frac{\sum_{x' \in \mathcal{X}(\hat{z})} p_r(g(\hat{f}(x')), x')}{\sum_{x' \in \mathcal{X}(\hat{z})} p(g(\hat{f}(x')), x')} & f(x') &= g(\hat{f}(x')) \\
 &= \frac{\sum_{x' \in \mathcal{X}(\hat{z})} p_r(g(\hat{z}), x')}{\sum_{x' \in \mathcal{X}(\hat{z})} p(g(\hat{z}), x')} & \hat{f}(x') &= \hat{z} \text{ for } x' \in \mathcal{X}(\hat{z}) \\
 &= \frac{\sum_{x' \in \mathcal{X}(\hat{z})} p_r(z, x')}{\sum_{x' \in \mathcal{X}(\hat{z})} p(z, x')} & z &= g(\hat{z}) \\
 &= \frac{\sum_{x' \in \mathcal{X}(\hat{z})} p_r(x' \mid z) p_r(z)}{\sum_{x' \in \mathcal{X}(\hat{z})} p(x \mid z) p(z)} & \\
 &= \frac{\sum_{x' \in \mathcal{X}(\hat{z})} p_r(x' \mid z) p_r(z)}{\sum_{x' \in \mathcal{X}(\hat{z})} (p_G q(x' \mid z) + (1 - p_G) p_r(x' \mid z)) p(z)} & \text{Fig. 2} \\
 &= \frac{\sum_{x' \in \mathcal{X}(\hat{z})} p_r(x' \mid z) p_r(z)}{\sum_{x' \in \mathcal{X}(\hat{z})} p_r(x' \mid z) p(z)} & q(x' \mid z) &= p_r(x' \mid z) \\
 &= \frac{p_r(z)}{p(z)}.
 \end{aligned}$$

1109

This concludes the proof.

1110

1111

A.5 DATASETS AND LABELING DETAILS

1112

1113

We describe the datasets used in our experiments, including their construction, preprocessing, and the labeling schemes relevant for evaluating semantic calibration under distributional shift. Note that the numbers reported in Table 1 for $|\mathcal{D}|$ are 80% of those reported for the training set below since 20% of the training data is kept as validation (and thus excluded in \mathcal{D} and \mathcal{D}_r). See Appendix A.7.2 for details about the exact semantic splits.

1118

1119

1120

1121

1122

Senator Tweets. We use the Senator Tweets dataset KristinCoyote (2025), where each instance is created by cleaning the tweet text of any links or emojis. The labels used are *conservative* and *liberal*, depending on the senator’s political affiliation. The final dataset’s size is 79.2k rows for training and 19.6k for testing.

1123

1124

1125

1126

1127

Political Podcasts. We use a subset of the Political Podcasts dataset nbandhi (2025). The dataset consists of segments from podcasts with political discussions. The existing labels which we use are *conservative* and *liberal*, depending on the context. The final dataset’s size is 11.4k rows for training and 1k for testing.

1128

1129

1130

1131

1132

1133

Political Bias. We use a subset of the Political Bias Corpus surajkarakulath (2025), where each instance is a news article, labeled as *left*, *center* or *right*. The final training dataset’s size is 9.59k and we evaluate on 2.3k samples.

1131

AG News. We use a subset of the AG News dataset Zhang et al. (2015), where each instance is created by concatenating the article’s *title* and *description*. This subset is limited to the four most common categories: *World*, *Sports*, *Business*, and *Sci/Tech*. It consists of 120,000 samples in the

1134 training set and 6,700 samples in the test set. The original four-class categorization is retained for our
 1135 experiments.
 1136

1137 **Gutenberg.** This dataset is derived from eight books sourced from the Project Gutenberg archive
 1138 Project Gutenberg (2025): *Pride and Prejudice*, *Frankenstein*, *The Great Gatsby*, *The Odyssey*,
 1139 *Moby Dick*, *Meditations*, *Oliver Twist*, and *War and Peace*. Each text is segmented into semantically
 1140 coherent spans of 50–100 words using `nltk.tokenize`, followed by post-processing (trimming
 1141 or merging) to enforce the desired range. Tokens matching non-linguistic patterns such as “–”, “_”,
 1142 and various brackets are removed. To ensure balance, we sample uniformly across books, resulting in
 1143 a dataset where each book contributes equally. The final dataset is split stochastically into 5,120 and
 1144 1,280 samples for training and testing, respectively. Phrases are labeled by their book of origin.
 1145

1146 **COCO.** We use a subset of the Microsoft COCO dataset Lin et al. (2014a), which contains images
 1147 annotated with objects from 80 predefined categories. From the original dataset, we select the first
 1148 68,800 images and generate captions using Florence Yuan et al. (2021). We further split the dataset
 1149 55,000 training samples and 13,800 test samples. Samples are assigned a binary label based on the
 1150 presence of objects typically associated with physical activity: *bicycle*, *frisbee*, *snowboard*, *sports*
 1151 *ball*, *kite*, *baseball bat*, *baseball glove*, *skateboard*, *surfboard*, *tennis racket*. Images with at least one
 1152 such object are assigned a label of 1, and 0 otherwise.
 1153

1154 **CIFAR-10.** The CIFAR-10 dataset Krizhevsky (2009) contains 60,000 images drawn from 10
 1155 distinct object types, uniformly distributed. We use the standard training and test splits of 50,000 and
 1156 10,000 images, respectively. We define the semantic shifts using the original 10 labels and caption
 1157 the images using Florence Yuan et al. (2021).
 1158

1159 **CIFAR-100.** CIFAR-100 Krizhevsky (2009) extends the structure of CIFAR-10 to 100 fine-grained
 1160 object types grouped under 20 supercategories. Each fine-grained type appears with uniform frequency
 1161 (1%), and coarse-level groupings appear at 5%. We use the canonical training/test split of 50,000 and
 1162 10,000 samples. To conduct distribution shifts we use coarse labels, whereas evaluation is performed
 1163 on the lower-level semantics of fine labels, which makes the task more challenging.
 1164

1165 **ImageNet.** We sample the first 127,000 images from the ILSVRC-2012 ImageNet dataset Deng
 1166 et al. (2009) and generate captions using Florence Yuan et al. (2021). The resulting set is partitioned
 1167 into training 101,600 training samples and 25,400 test samples. To introduce a higher-level semantic
 1168 structure, we map the original 673 fine-grained labels to a custom taxonomy of 20 coarse categories
 1169 (see Appendix A.7.2 for details). These coarse labels are used to induce structured semantic variation
 1170 and evaluate calibration behavior under abstraction.
 1171

1172 **Hateful Memes.** We use the full Hateful Memes dataset (Kiela et al., 2021) dataset. It comes with
 1173 existing labels (hateful or non-hateful). We use Florence to caption the images and concatenate it
 1174 with the provided text. The training dataset size is 6.8k and the test dataset size is 1.7k.
 1175

1176 **Activity Net.** We use the full Activity Net dataset (Caba Heilbron et al., 2015) dataset. We map all
 1177 `nodeName` labels to their `parentName` and obtain a training dataset of size 11.9k with 52 different
 1178 coarse labels. The captions are taken directly from (Krishna et al., 2017).
 1179

1180 **UrbanSound8k.** We use the full UrbanSound8k dataset Salamon et al. (2014), which consists of
 1181 audio recordings from 10 acoustic event categories (e.g., *dog bark*, *children playing*, *gun shot*, etc.).
 1182 As the dataset does not include an official train/test split, we partition it ourselves using an 80/20
 1183 random split. Each audio sample is captioned using Qwen-Audio Chu et al. (2023) with the prompt:
 1184 *“Describe this audio.”* The resulting captions are used to evaluate semantic calibration in the auditory
 1185 modality.
 1186

1187 **VGG Sound.** We construct this dataset from the original VGG Sound collection Chen et al. (2020)
 1188 by mapping its 309 fine-grained acoustic categories to a custom taxonomy of 14 coarse labels (see
 1189 Appendix A.7.2 for details). To avoid redundant computation, we reuse the natural language captions
 1190 provided in Mei et al. (2023), specifically the model’s *answer-1* to *question-1* for each sample.
 1191

1188 A.6 CAPTIONING MODEL PROMPTS.
 1189

- **Text Modality:** no prompt is necessary since no captioning model is used. Text samples are used directly as input.
- **Image Modality:** we pass the standard prompt: <MORE_DETAILED_CAPTION> to Florence Yuan et al. (2021).
- **Audio Modality:** we pass the prompt: "Describe the sounds in this clip" for VGGSound and "Describe this audio" for UrbanSound8k.
- **Video Modality:** we use the generated captions from the "Dense-Captioning Events in Videos" paper (Krishna et al., 2017) .

1199 A.7 REPRODUCIBILITY

1200 A.7.1 CODE AND IMPLEMENTATION

1203 All code necessary to reproduce our experiments is available at <https://github.com/BeyondDeepFakeDetection/Beyond-Deep-Fake-Detection>.

1205 The repository includes:

- The scripts that were used for image and audio captioning, and to semantically split books into phrases.
- The script used to generate distorted distribution datasets given the general dataset.
- The LLM fine-tuning script, with all hyperparameters matching those reported in the main text.
- The JSON files with custom low- to high-level label mappings, that were used in the dataset engineering part.
- The script used to compute the per-token and total probability of a text under a fine-tuned LLM.
- A standalone implementation of our *semantic calibration* pipeline, with:
 - Rebalancing of the probability mass using ρ .
 - Online filtering mechanism using rejection sampling.
 - Performance metrics computation (e.g. Deception Reduction).

1223 A.7.2 DATASETS AND FINE-TUNED MODELS

1224 All datasets and fine-tuned models utilized in our experiments are publicly accessible via our Hugging Face repository:

1227 <https://huggingface.co/BeyondDeepFakeDetection>.

1229 Each dataset includes: general training set, general test set, baseline split training set, mild split 1230 training set, moderate split training set, and severe split training set. Each fine-tuned model has been 1231 trained on its corresponding training split.

1233 A.8 COMPUTE RESOURCES

1235 We detail the compute requirements for each component of our pipeline to support reproducibility.

- **LLM fine-tuning (GPT-2 Small):** Fine-tuning GPT-2 (124M parameters) on a dataset of \sim 50k samples such as COCO no sports, for 5 epochs, was completed in less than 2 hours on a single NVIDIA RTX 2080 GPU.
- **Florence captioning model inference:** Caption generation using Microsoft's large Florence model was performed on an NVIDIA RTX A6000 GPU (48GB VRAM). Processing a dataset of 10k images requires approximately 2.5 hours.

- **Qwen-Audio captioning inference:** Audio captioning with Qwen-Audio was run on an NVIDIA RTX A6000. Captioning the UrbanSound8k dataset (8.6k rows) required approximately 2.5 hours.
- **LLM inference (per-token log probabilities):** For each test dataset, we computed token-wise log probabilities using a fine-tuned GPT-2 model. Inference over the evaluation set took less than ~ 15 minutes even on our largest datasets ($> 35k$ rows) using a single NVIDIA RTX 2080 GPU.
- **Rejection sampling process:** Our semantic calibration mechanism performs 1000 independent rejection sampling runs to compute statistics (mean, std) for distribution alignment. This process is CPU-only, completes in under 5 minutes even for the largest datasets and has negligible compute cost.

A.9 SEMANTIC CALIBRATION VERSUS DEEPFAKE DETECTION

We illustrate the difference between semantic calibration and deepfake detection in Fig. 7.

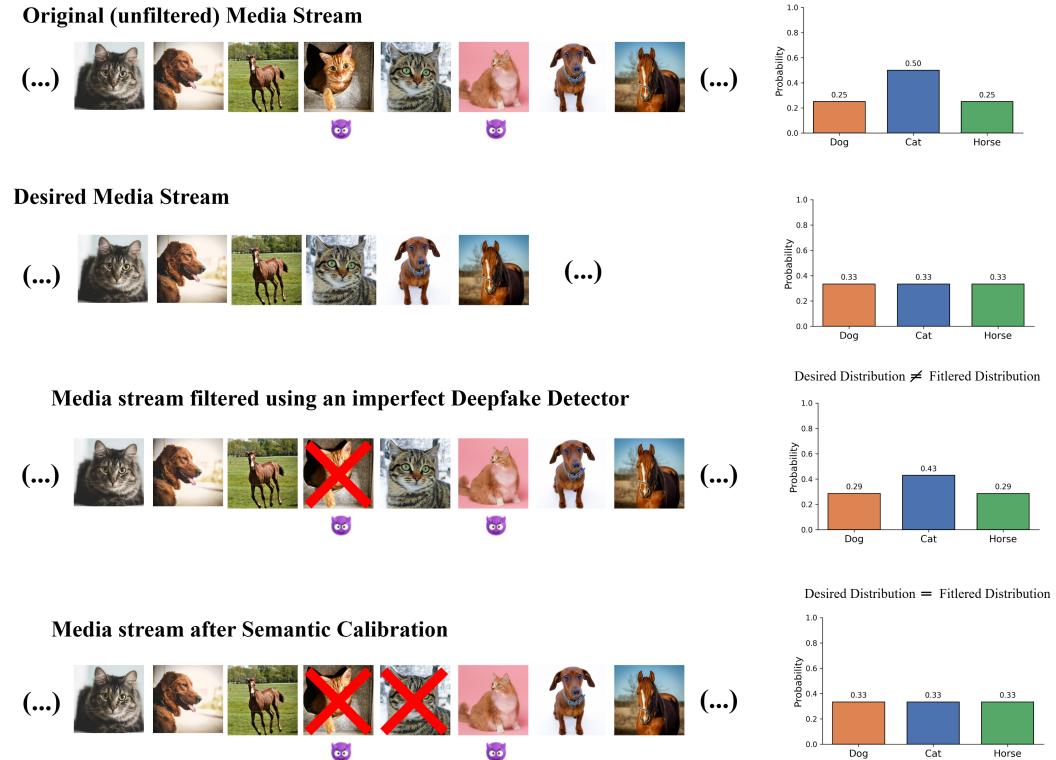


Figure 7: We illustrate an unfiltered media stream containing images of dogs, cats, and horses. The stream induces a non-uniform distribution p , as 2 of the 4 cat images (50%) are deepfakes (marked with the evil emoji). Suppose the deepfake detector is imperfect, as Theorem 1 suggests any detector may eventually be, and removes only 1 of the 2 fake cat images. The filtered distribution therefore remains skewed relative to the desired target. In contrast, applying our semantic calibration framework in the latent semantic space correctly adjusts the distribution p_r by removing two cat instances regardless of whether they are fake or real. This example shows that our method is not a conventional deepfake detector: even when detection fails, semantic calibration rebalances manipulated media streams toward the target distribution.

1296
1297

A.10 ADDITIONAL FIGURES FOR CALIBRATION PERFORMANCE

1298
1299
1300
1301

In Fig. 8 and Fig. 9, we present two additional figures that complement Fig. 4 from the main text. These provide further insight into the behavior of the distributions under different dataset shift scenarios. To further illustrate the dynamics of our method in an online setting, we provide a live animation available in the GitHub repository listed in Appendix A.7.1.

1302

1303

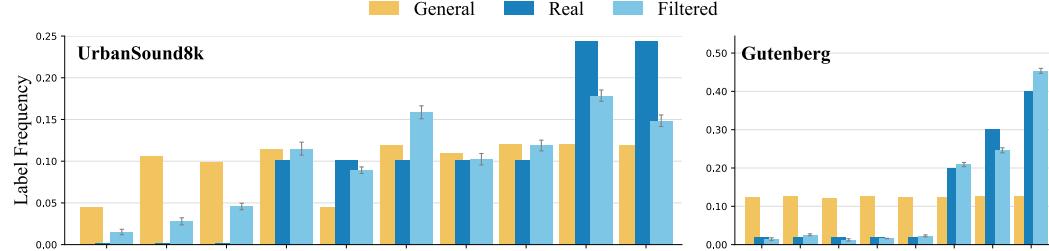
1311
1312
1313

Figure 8: Comparison of the **filtered distribution** $p^\phi(z)$, the **real distribution** $p_r(z)$, and the **general distribution** $p(z)$ under a *severe shift* scenario. Results are shown for the UrbanSound8K and Gutenberg datasets. Class labels on the horizontal axis are ordered by increasing frequency in $p_r(z)$.

1314

1315

1316

1317

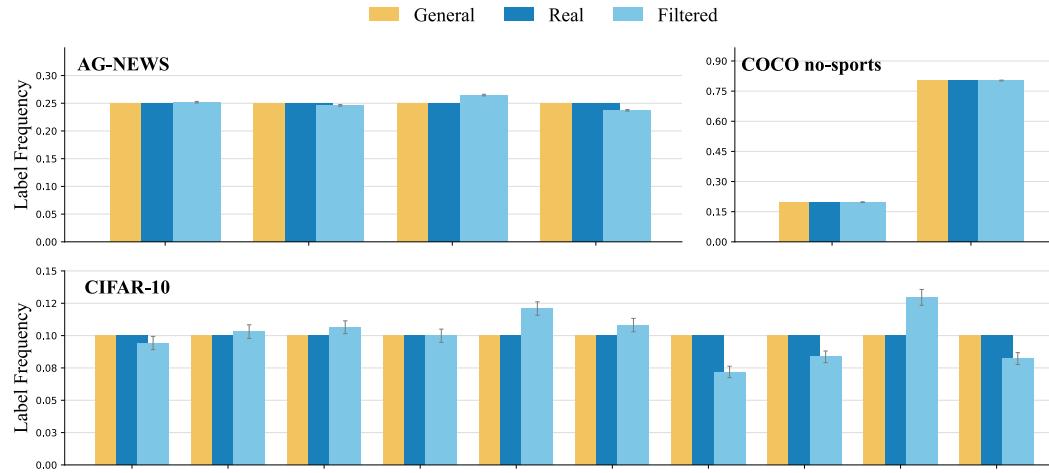
1325
1326
1327
1328
1329
1330
1331
1332
1333
1334

Figure 9: Visualization of the **filtered distribution** $p^\phi(z)$, the **real distribution** $p_r(z)$, and the **general distribution** $p(z)$ under a *baseline shift* scenario. Shown are results for AG-NEWS, COCO (no-sports), and CIFAR-10. Class labels along the horizontal axis are ordered by increasing frequency under $p_r(z)$. Our approach introduces minimal distortion, as expected for baseline scenarios.

1335

1336

1337

A.11 ADDITIONAL FIGURES FOR EXPLAINABILITY

1338

1339

1340

To complement Fig. 1 and Fig. 5, we show additional saliency maps in Fig. 10, Fig. 12 and Fig. 14 with their corresponding rolling acceptance probabilities in Fig. 11, Fig. 13 and Fig. 15, respectively.

1341

1342

1343

1344

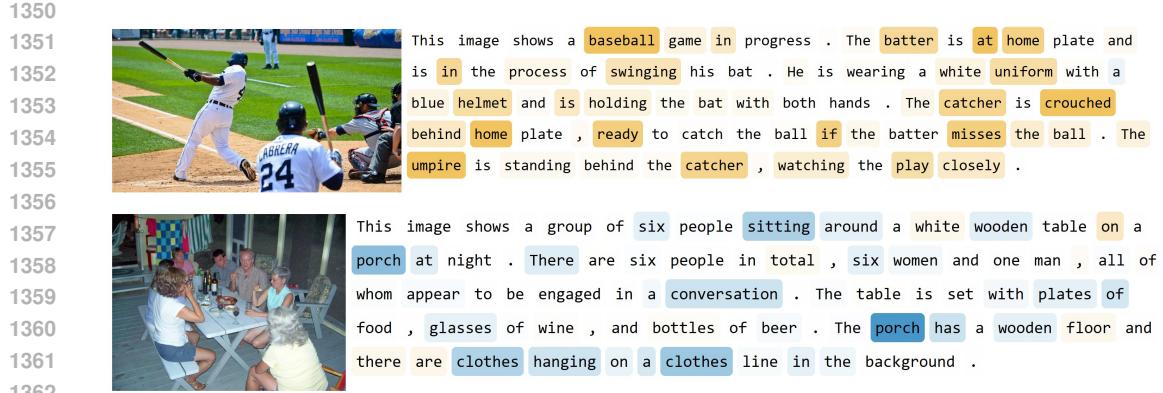
1345

1346

1347

1348

1349



1363 Figure 10: Saliency map showcasing the explainability of semantic calibration on two test images
1364 from the COCO dataset. Tokens highlighted in **blue favor acceptance** ($\Delta_i > 0$) while those
1365 highlighted in **orange favor rejection** ($\Delta_i < 0$). As expected, words highlighted in orange intuitively
1366 favor rejection (e.g., baseball, batter, catcher, umpire) and those in blue favor acceptance
1367 (e.g. sitting, conversation, porch).

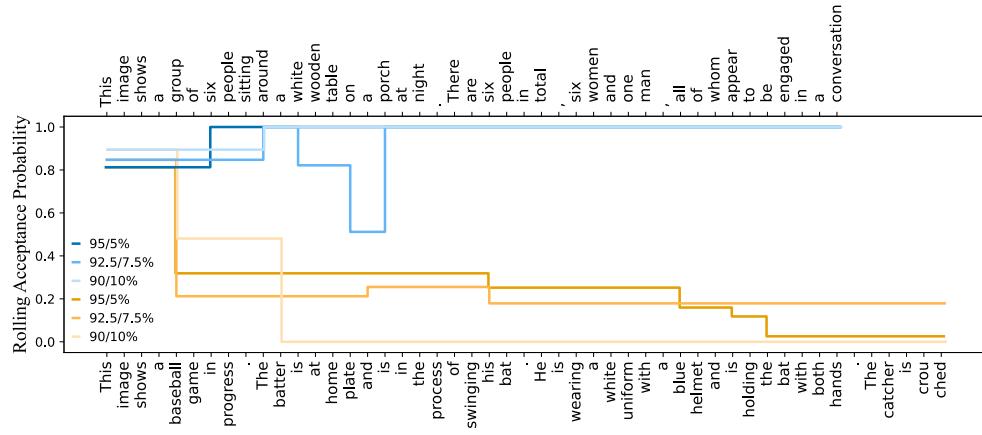
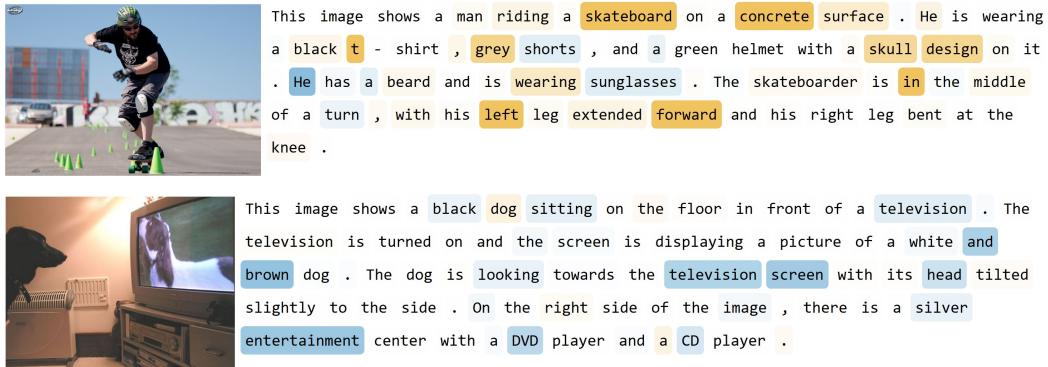


Figure 11: Rolling acceptance probabilities for the two images shown in Fig. 10



1399 Figure 12: Saliency map showcasing the explainability of semantic calibration on two test images from
1400 the COCO dataset. Tokens highlighted in **blue favor acceptance** ($\Delta_i > 0$) while those highlighted
1401 in **orange favor rejection** ($\Delta_i < 0$). As expected, words highlighted in orange intuitively favor
1402 rejection (e.g., skateboard, forward) and those in blue favor acceptance (e.g. television,
1403 screen, entertainment).

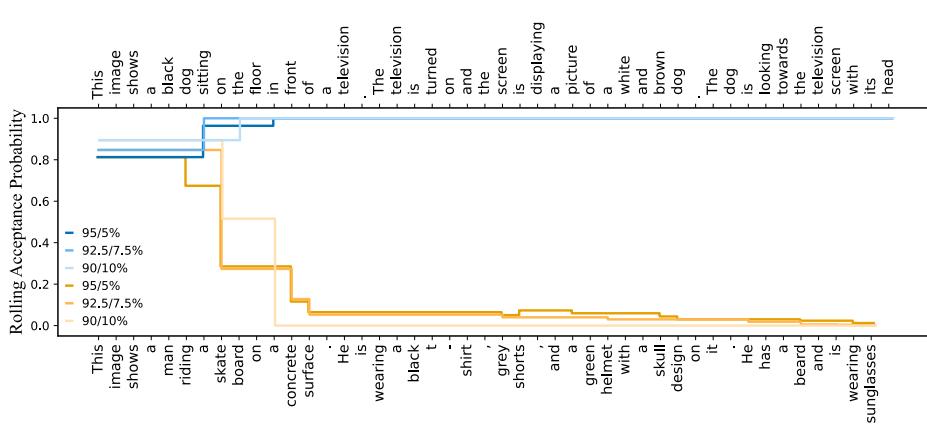


Figure 13: Rolling acceptance probabilities for the two images shown in Fig. 12

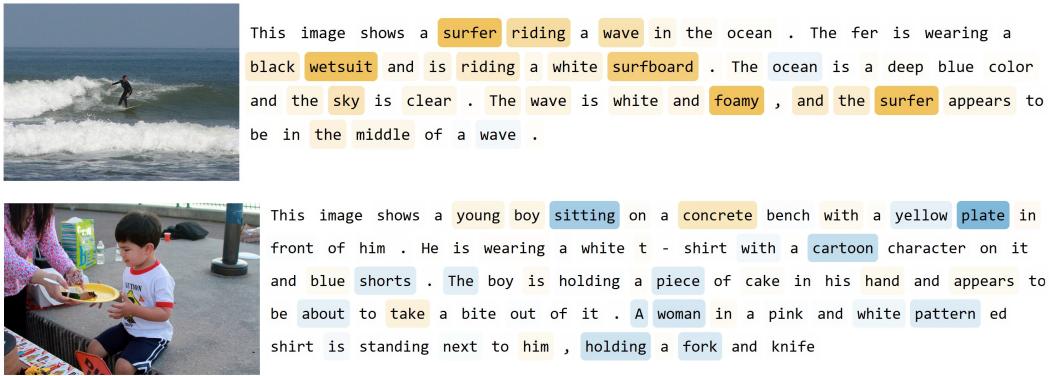
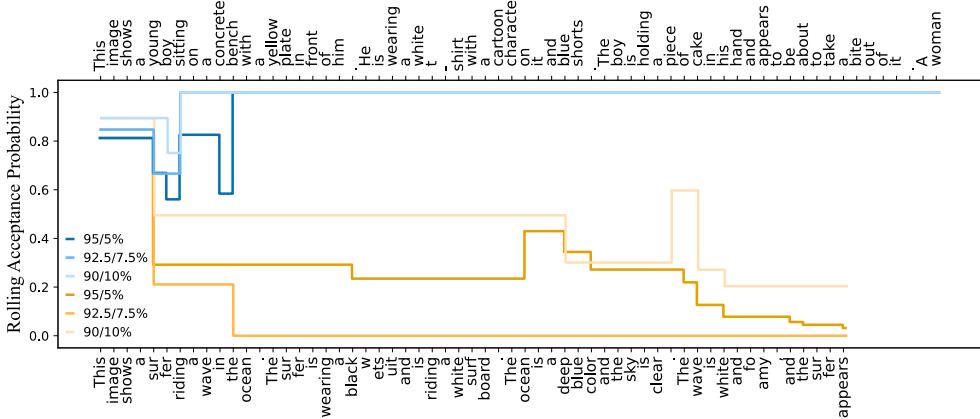
Figure 14: Saliency map showcasing the explainability of semantic calibration on two test images from the COCO dataset. Tokens highlighted in **blue favor acceptance** ($\Delta_i > 0$) while those highlighted in **orange favor rejection** ($\Delta_i < 0$). As expected, words highlighted in orange intuitively favor rejection (e.g., surfer, wetsuit, surfboard) and those in blue favor acceptance (e.g., sitting, plate, fork).

Figure 15: Rolling acceptance probabilities for the two images shown in Fig. 14