

Reinforcement Learning-Based Prosody Optimization and Adaptive Singing Style Modeling

Ningcai Shen

Department of Electrical Engineering and Computer Science, University of Kansas, Lawrence, KS, USA.

ning1982@ku.edu

Ronald L. Bush

Department of Computer Science, Colorado State University, Fort Collins, CO, USA.

ronaldbush100@colostate.edu

Abstract

Singing voice synthesis has advanced rapidly through neural sequence modeling and large-scale audio corpora, but expressive control over prosody and singing style remains an open systems challenge. This paper presents a long-form examination of reinforcement learning-based prosody optimization and adaptive singing style modeling. The central argument is that prosody should be treated not as a fixed set of acoustic contours but as an interactive policy over expressive decisions that must be learned, constrained, and governed at the system level. The paper develops a reference architecture in which a prosodic policy network interacts with a differentiable singing synthesizer and a reward model that encodes musicality, intelligibility, singer identity, and stylistic appropriateness. The discussion analyzes structural trade-offs among reward sparsity, credit assignment, exploration, training stability, deployment latency, and environmental sustainability. It also examines style encoding, reference provenance, and identity governance. By comparing speech prosody transfer, generative music models, and singing voice synthesis, the paper identifies conditions under which reinforcement learning can improve adaptive style modeling without destabilizing the acoustic backbone. The analysis emphasizes that technical gains must be evaluated alongside sociotechnical safeguards, including dataset consent, cultural diversity, fairness, privacy, and equitable access. The contribution is a critical synthesis of architectural, infrastructural, and policy considerations for building controllable, adaptive, and responsibly governed singing voice systems.

Keywords

reinforcement learning, prosody optimization, singing voice synthesis, adaptive style modeling, expressive speech synthesis, sociotechnical governance.

1. Introduction

Modern expressive singing synthesis sits at the intersection of reinforcement learning, sequence modeling, and music information retrieval. The formal treatment of sequential decision making under uncertainty [1] and its deep learning extensions [2] provide conceptual tools for understanding prosodic variation as a policy over expressive actions rather than a static post-processing stage. Attention-based sequence modeling [3] has also reshaped neural speech and audio architectures, enabling systems to condition generation on long-range context and structured musical scores. This conceptual shift is important because singing

requires continuous control over pitch, timing, energy, timbre, vibrato, and articulation, all of which vary with genre, language, singer physiology, and emotional intention.

High-quality speech synthesis systems have demonstrated that phonetic and duration prediction can be separated from vocoding to produce natural speech [4], [5]. These systems provide a useful technical foundation for singing synthesis, but they do not fully address the evaluative and interactive nature of expressive performance. In speech synthesis, intelligibility and naturalness are primary objectives. In singing synthesis, musicality, stylistic identity, ornamentation, and emotional plausibility are equally important. These qualities are difficult to specify as simple regression targets because they involve aesthetic judgments that vary across listeners, cultures, and performance traditions. A reinforcement learning formulation is therefore appealing: it can optimize a policy against a learned or explicitly defined reward signal while preserving the flexibility needed for stylistic adaptation.

Style modeling approaches such as unsupervised style tokens [6] show that global and local expressive factors can be learned from speech data. However, singing voice synthesis requires more structured adaptation because style is not monolithic. It includes singer identity, genre conventions, ornamentation patterns, regional performance practice, and momentary expressive decisions. A system-level discussion must therefore consider how these dimensions can be represented, transferred, and governed. The next sections develop this discussion by connecting related work, architectural design, reinforcement learning formulation, infrastructure, and policy implications.

2. Related Work and System Context

The emergence of large-scale, richly annotated singing corpora has made it possible to train systems across multiple languages, techniques, and score-aligned conditions [7]. Such corpora reduce the mismatch between controlled laboratory speech data and real-world singing performance by providing aligned musical scores, phoneme sequences, and audio. At the same time, they introduce questions about annotation reliability, singer privacy, cultural representation, and consent. The value of a large corpus is not only its size but also the quality and governance of its metadata. A corpus that is large but stylistically narrow may still produce models that perform poorly on underrepresented singing traditions.

Prosody transfer in speech synthesis has often been formulated as a conditional reconstruction problem in which a reference audio or learned style embedding guides the decoder [8], [9]. These methods are effective for transferring global contours such as speaking rate and pitch range, but they do not close the loop through explicit musical reward signals. They also tend to treat prosody as a latent condition rather than an interactive policy that can be adjusted under multiple objectives. When such methods are extended to singing, they may produce plausible audio that nevertheless lacks the deliberate expressive choices associated with expert performance.

Reinforcement learning has been applied to sequence generation through conservative policy updates and task-based rewards [10], [11]. These approaches address a central dilemma in generative modeling: fine-tuning a sequence model with a scalar reward can produce degenerate or unstable outputs if the reward is misspecified. In singing synthesis, the same issue arises when a reward model places excessive weight on pitch accuracy and thereby suppresses expressive ornamentation. Conservative policy updates, reference-based regularization, and reward shaping are therefore important design tools.

Waveform-level generative models [12] enabled high-fidelity audio synthesis, while latent variable methods [13] and adversarial training [14] provided mechanisms for controlling variability and realism. In parallel, symbolic music modeling [15] showed that factorized representations of pitch, timing, and dynamics can improve long-range coherence. Singing synthesis integrates these strands because it requires both a musical score interpretation and a waveform or spectral audio generation pipeline. The relevant research therefore spans speech, music, and general audio generation, and system designers must reconcile the different assumptions embedded in each lineage.

3. Architectural Foundations and System-Level Design

At the system level, a reinforcement learning-based singing synthesizer can be organized into three interacting subsystems: a score and lyric encoder, a prosodic policy network, and a differentiable acoustic decoder. The score and lyric encoder transforms musical notation and phoneme sequences into structured contextual representations. The prosodic policy network selects actions that modify pitch trajectories, durations, energy contours, vibrato depth, and articulation. The acoustic decoder then produces a spectral representation or latent audio that can be rendered by a neural vocoder. This decomposition is useful because it isolates the expressive decision layer from the high-dimensional audio generation layer.

The prosodic policy network may observe past acoustic context and score features, then emit modifications that are applied by the decoder. The action space can be continuous, discrete, or mixed. Continuous actions allow fine-grained control of pitch curves and dynamics, while discrete actions may select among learned ornamentation primitives. A mixed action space often emerges in practice because phoneme-level decisions interact with frame-level acoustic trajectories. For example, a discrete choice of a melisma pattern may be followed by continuous adjustments to its timing and amplitude.

An important architectural trade-off concerns the placement of the policy relative to the decoder. If the policy operates on low-dimensional prosodic parameters, optimization is relatively stable, but some expressiveness may be lost because the decoder cannot adjust to novel combinations. If the policy directly manipulates latent decoder states, expressiveness can increase, but credit assignment becomes more difficult and training may become unstable. A layered architecture often balances these concerns by using high-level style commands, mid-level prosodic parameters, and low-level acoustic residuals. Governance and interpretability benefit when the high-level commands are human-readable and can be inspected or modified by a human musician.

The choice of decoder also affects the system. A deterministic decoder may be easy to optimize but may lack the variation needed for expressive performance. A stochastic decoder can produce multiple plausible realizations of the same score, which is essential for rehearsal-like exploration. However, stochastic generation complicates reward evaluation because the same action can lead to different acoustic outcomes. System designers often address this by using multiple samples per update and by averaging reward estimates. This increases computational cost but improves the stability of policy learning.

4. Prosody Optimization as a Reinforcement Learning Problem

Generative audio models have demonstrated that text or symbolic conditions can yield convincing musical output. Systems such as Jukebox [16], MusicLM [17], MusicGen [18], and AudioLM [19] differ in their representational choices, but all use large-scale pretraining to learn acoustic regularities. These models provide a useful diagnostic reference for singing

synthesis because they separate high-level conditioning from waveform generation. However, they do not usually treat prosody as an explicitly optimized control policy with a reward signal that reflects musicality, intelligibility, and style consistency. Singing systems can learn from these models while adding a more explicit decision-making layer.

Human preference learning provides one way to specify a reward that is otherwise hard to write as a deterministic objective [20]. In this setting, a reward model is trained from comparative judgments, and the prosodic policy is optimized against that model. This arrangement can capture subtle judgments such as the appropriateness of a melisma or the naturalness of a vibrato onset. It also introduces governance obligations because preference data may encode idiosyncratic, cultural, or even biased notions of good singing. If the preference model is trained primarily on listeners from one musical background, it may penalize vocal timbres or ornamentation practices from other traditions.

The optimization problem can be framed as a sequential decision process in which the policy receives a state summarizing score position, phoneme context, singer identity, and recent acoustic history. The policy emits a prosodic modification, and the environment returns an acoustic observation and a reward. The objective is to learn a policy that maximizes the expected cumulative reward over a complete phrase while remaining close to a pretrained reference policy. Closeness to a reference is important for stability because it prevents large policy updates from producing unnatural phonation, unstable pitch, or loss of lyric intelligibility.

Reward shaping is critical in this process. A monolithic quality reward may be too sparse to guide learning early in training. Curriculum rewards can provide denser feedback by first encouraging pitch accuracy and phonetic intelligibility, then adding style matching and emotional appropriateness. But reward shaping is not neutral. It embeds assumptions about which musical features matter and how they should be prioritized. Therefore, reward design should be treated as a governance artifact rather than a purely technical convenience. Documenting reward components, their weights, and their intended behavioral effects is necessary for auditability.

Exploration is another structural challenge. If the policy is initialized from a supervised model, it may already produce acceptable performances and therefore have little incentive to explore unusual but valuable stylistic variations. Entropy bonuses, noise schedules, or a population of diverse policies can encourage exploration, but each approach introduces instability. A population-based approach also raises deployment and evaluation costs because multiple policies must be maintained. The system-level decision is not merely which exploration method to use but how much expressive variance is desirable for a given application context.

5. Adaptive Singing Style Modeling and Style Governance

Adaptive style modeling requires the system to infer, store, and transfer expressive attributes across performances. Style tokens [6] provide an unsupervised basis for representing global style, but singing requires more structured adaptation because genre, ornamentation, and singer physiology interact. A reference encoder may extract style from a short audio clip, while the policy conditions on that embedding. The challenge is to avoid entangling singer identity with musical interpretation. Identity should be stable and protected, whereas interpretation should be flexible and controllable. If identity and interpretation are entangled, the system may inadvertently transfer a singer identity when only a performance style was intended.

System-level style governance involves maintaining provenance of style references. If a singer identity can be cloned from audio, the system must distinguish between authorized identity modeling and unauthorized voice replication. Style transfer may also raise cultural property questions when ornamentation is associated with particular traditions. An architecture that records style provenance and provides granular control can support responsible deployment. For example, a user might be allowed to transfer a general expressive feature such as legato phrasing while being prevented from cloning a specific singer identity.

Large singing corpora with multimodal annotations support fine-grained style conditioning [7]. However, the size of a corpus does not automatically imply representational fairness. A global corpus may overrepresent commercial genres and underrepresent regional, non-Western, or historical singing practices. The adaptive model should therefore include evaluation protocols that measure style fidelity across multiple traditions, not only aggregate acoustic similarity. This requires collaboration with musicologists, performers, and community representatives who can identify culturally meaningful features that automatic metrics may overlook.

Style adaptation also has an interactive dimension. A performer or producer may wish to adjust style in real time by providing feedback such as more legato, less vibrato, or brighter timbre. Reinforcement learning can support this interaction by treating user feedback as reward, but the system must distinguish between stable preferences and momentary adjustments. If the policy adapts too quickly, it may overfit to a single user and lose generalization. If it adapts too slowly, it may feel unresponsive. Architectural mechanisms such as online adaptation with bounded parameter changes can help maintain this balance.

6. Training Infrastructure, Deployment, and Sustainability

Training a prosodic policy against a neural vocoder is computationally demanding. The system must generate many candidate performances, evaluate them with a reward model, and update both the policy and possibly the decoder. Multi-stage training is often necessary. First, the acoustic backbone is pretrained on supervised singing data. Second, a reward model is trained on preference judgments. Third, the prosodic policy is fine-tuned with reinforcement learning. This decomposition improves reproducibility and allows intermediate auditing, but it also creates dependencies. If the acoustic backbone changes, the reward model and policy may need to be retrained or recalibrated.

Deployment latency is another structural constraint. A policy that requires many iterative refinement steps may not be suitable for real-time composition or live performance. Model distillation, quantization, and fixed inference budgets are common responses, but these techniques can reduce the very expressiveness that the policy was designed to learn. System designers therefore need to choose between offline rendering with rich optimization and online rendering with constrained policy capacity. Hybrid deployments may use offline optimization for studio production and a distilled policy for interactive preview. Each mode should be evaluated according to its intended use rather than a single latency benchmark.

Sustainability considerations are increasingly relevant. Large-scale audio models require substantial energy for training and inference. A reinforcement learning pipeline that samples many candidate performances increases this footprint. One design response is to optimize in a compact prosodic space before decoding only the most promising candidates. Another response is to reuse pretrained backbones across multiple languages and styles, reducing the

need for independent retraining. Governance frameworks should report energy use and require a proportionality test between expressive benefit and environmental cost. This is especially important when the system is deployed at scale in consumer applications where many users may trigger computationally intensive inference.

The training infrastructure must also support reproducibility and debugging. Reinforcement learning for audio generation can fail in subtle ways, such as gradually reducing vocal diversity while still receiving high reward scores. Logging and visualization tools should track not only reward magnitude but also the distribution of generated prosodic features over time. If the policy collapses to a narrow range of performances, the system should detect this shift and trigger human review. This kind of infrastructure is not merely an engineering convenience; it is necessary for responsible oversight.

7. Robustness, Fairness, and Policy Implications

Algorithmic accountability research has identified recurring ethical problems in data-driven systems, including opacity, unfairness, and the risk of encoding contested values [21]. In singing voice synthesis, these concerns manifest in several ways. A reward model trained on narrow listener preferences may penalize vocal timbres associated with marginalized groups. A style encoder may infer sensitive attributes from singing audio and use them in ways that are not transparent to users. Governance mechanisms must therefore include external audits and disaggregated evaluation. Audits should examine how the system performs across languages, genders, vocal ranges, and singing traditions, and should not be limited to average performance metrics.

Sociotechnical fairness research emphasizes that algorithmic outcomes cannot be assessed in isolation from the institutions and practices in which they are embedded [22]. A singing synthesis system is not merely a model; it is embedded in music production platforms, copyright regimes, fan communities, and cultural economies. Fairness analysis must consider who can contribute training data, who is represented in style conditioning, who can access the model, and who bears the risk of misuse. Technical metrics such as mean opinion score or pitch accuracy do not capture these distributional questions. A system that performs well on average may nevertheless exclude minority singing practices or concentrate benefits among already dominant commercial genres.

Large pretrained systems can reproduce patterns from their training data in ways that appear fluent but may be socially harmful [23]. In singing synthesis, the equivalent risk is the generation of performances that imitate a protected singer identity or that blend cultural styles without attribution. Policy responses include consent-based data collection, singer identity verification, style provenance documentation, and watermarking of generated audio. Such mechanisms should be designed before deployment, not retrofitted after harm has occurred. Watermarking can help identify synthetic audio, but it must be implemented without degrading the musical quality or becoming trivially removable.

Regulatory alignment is also important. Voice synthesis is becoming subject to privacy and biometric regulations in several jurisdictions. A reinforcement learning system that updates a singer identity embedding from user feedback may inadvertently create a model that is subject to data protection obligations. System architects should treat the policy, reward model, and style encoder as potentially personal data processing components and apply privacy-preserving learning where feasible. This may include differential privacy during reward model training, data minimization in style reference storage, and clear user consent for

identity-related adaptation. The design should be able to demonstrate compliance through technical documentation and institutional review.

8. Evaluation and Cross-Domain Comparison

Evaluation of prosody optimization cannot rely on a single scalar metric. Intonation accuracy, lyric intelligibility, style fidelity, emotional plausibility, and singer consistency are partially independent objectives. Human listening tests remain central, but they are costly and can be biased by rater demographics and instructions. Automated metrics may capture pitch and timing error but often fail to detect culturally meaningful ornamentation. A robust evaluation framework should combine objective acoustic analysis, preference-model accuracy, and structured human assessment. The framework should also report confidence intervals and subgroups so that performance variation is visible rather than hidden in aggregate scores.

Cross-domain comparison with speech prosody transfer and generative music models is instructive. Speech systems prioritize intelligibility and naturalness under tight latency constraints, whereas generative music systems prioritize long-range structure and stylistic coherence. Singing synthesis inherits both sets of tensions: it must preserve linguistic content while serving a musical and emotional goal. Reinforcement learning is useful precisely because it can handle multiple objectives, but the priority ordering among those objectives is a normative choice. System designers should make that ordering explicit and configurable. A production system may allow different reward weight profiles for choral, popular, operatic, or traditional singing.

A case illustration can clarify these trade-offs. Suppose a system is asked to generate a traditional folk song with ornamentation from a particular regional style. A purely supervised model may produce a smoothed and standardized performance because the training corpus contains few examples. A reinforcement learning policy can be rewarded for matching annotated ornamentation events, but only if the corpus includes those annotations and the reward model is able to recognize them. If the reward model is trained on listeners unfamiliar with the tradition, it may inadvertently suppress exactly those features. This shows that evaluation and reward design are deeply tied to cultural competence. It also suggests that evaluation panels should include listeners with relevant stylistic knowledge.

The comparison with general audio generation also highlights the importance of control granularity. Large generative music models can create impressive long-form music, but users often struggle to make precise adjustments to a single phrase. Singing synthesis applications may require fine-grained control over a single note, syllable, or ornamentation. Reinforcement learning can provide this control by framing user feedback as part of the reward signal. However, the interface must be designed so that users understand what the system is optimizing and can intervene when the policy produces an unintended result. Instrumentation and logging should record not only what the user requested but also how the policy interpreted that request.

9. Conclusion

Reinforcement learning-based prosody optimization offers a promising path toward adaptive, controllable singing voice synthesis, but it also introduces structural complexity that must be managed at the system level. This paper has argued that prosody should be treated as an interactive policy over expressive decisions, supported by a pretrained acoustic backbone, a preference-informed reward model, and a governance layer that constrains style transfer and identity use. The architectural design must balance expressiveness against stability, latency

against refinement, and optimization depth against sustainability. These trade-offs cannot be resolved by any single algorithmic choice; they require deliberate system design and ongoing institutional oversight.

Future work should develop culturally grounded reward models, auditable style provenance, and evaluation protocols that reflect the diversity of singing traditions. The integration of reinforcement learning into singing synthesis should not proceed as a purely technical optimization project. It must engage with musicology, ethics, and platform policy to ensure that expressive technologies are both innovative and accountable. Progress in this area will require not only advances in machine learning but also sustained interdisciplinary collaboration and the development of governance practices that can keep pace with rapid technical change.

References

1. Sutton, R. S., & Barto, A. G. (2018). Reinforcement learning: An introduction (2nd ed.). MIT Press.
2. Mnih, V., Kavukcuoglu, K., Silver, D., Rusu, A. A., Veness, J., Bellemare, M. G., Graves, A., Riedmiller, M., Fidjeland, A. K., Ostrovski, G., Petersen, S., Beattie, C., Sadik, A., Antonoglou, I., King, H., Kumaran, D., Wierstra, D., Legg, S., & Hassabis, D. (2015). Human-level control through deep reinforcement learning. *Nature*, 518(7540), 529-533.
3. Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, Ł., & Polosukhin, I. (2017). Attention is all you need. *Advances in Neural Information Processing Systems*, 30, 5998-6008.
4. Shen, J., Pang, R., Weiss, R. J., Schuster, M., Jaitly, N., Yang, Z., Chen, Z., Zhang, Y., Wang, Y., Skerry-Ryan, R., Saurous, R. A., Agiomyrgiannakis, Y., & Wu, Y. (2018). Natural TTS synthesis by conditioning WaveNet on mel spectrogram predictions. 2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP), 4779-4783.
5. Ren, Y., Ruan, Y., Tan, X., Qin, T., Zhao, S., Zhao, Z., & Liu, T.-Y. (2019). FastSpeech: Fast, robust and controllable text to speech. *Advances in Neural Information Processing Systems*, 32, 3171-3180.
6. Wang, Y., Stanton, D., Zhang, Y., Skerry-Ryan, R. J., Battenberg, E., Shor, J., Xiao, Y., Ren, F., Jia, Y., & Saurous, R. A. (2018). Style tokens: Unsupervised style modeling, control and transfer in end-to-end speech synthesis. *Proceedings of the 35th International Conference on Machine Learning*, 5180-5189.
7. Zhang, Y., Pan, C., Guo, W., Li, R., Zhu, Z., Wang, J., ... & Zhao, Z. (2024). Gtsinger: A global multi-technique singing corpus with realistic music scores for all singing tasks. *Advances in Neural Information Processing Systems*, 37, 1117-1140.
8. Skerry-Ryan, R. J., Battenberg, E., Xiao, Y., Wang, Y., Stanton, D., Shor, J., Weiss, R. J., Clark, R., & Saurous, R. A. (2018). Towards end-to-end prosody transfer for expressive speech synthesis with Tacotron. *Proceedings of the 35th International Conference on Machine Learning*, 4693-4702.
9. Klimkov, V., Ronanki, S., Rohnke, J., & Merritt, T. (2019). Fine-grained robust prosody transfer for single-speaker neural text-to-speech. *Proceedings of Interspeech 2019*, 4430-4434.

10. Jaques, N., Gu, S., Turner, R. E., & Eck, D. (2017). Sequence tutor: Conservative fine-tuning of sequence generation models with KL-control. *Proceedings of the 34th International Conference on Machine Learning*, 1645-1654.
11. Schulman, J., Wolski, F., Dhariwal, P., Radford, A., & Klimov, O. (2017). Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*.
12. van den Oord, A., Dieleman, S., Zen, H., Simonyan, K., Vinyals, O., Graves, A., Kalchbrenner, N., Senior, A., & Kavukcuoglu, K. (2016). WaveNet: A generative model for raw audio. *arXiv preprint arXiv:1609.03499*.
13. Kingma, D. P., & Welling, M. (2014). Auto-encoding variational Bayes. *Proceedings of the 2nd International Conference on Learning Representations*.
14. Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., & Bengio, Y. (2014). Generative adversarial nets. *Advances in Neural Information Processing Systems*, 27, 2672-2680.
15. Hawthorne, C., Stasyuk, A., Roberts, A., Simon, I., Huang, C.-Z. A., Dieleman, S., Elsen, E., Engel, J., & Eck, D. (2019). Enabling factorized piano music modeling and generation with the MAESTRO dataset. *Proceedings of the 7th International Conference on Learning Representations*.
16. Dhariwal, P., Jun, H., Payne, C., Kim, J., Radford, A., & Sutskever, I. (2020). Jukebox: A generative model for music. *arXiv preprint arXiv:2005.00341*.
17. Agostinelli, A., Denk, T. I., Borsos, Z., Engel, J., Verzetti, M., Caillon, A., Huang, Q., Jansen, A., Roberts, A., TagliasTagliasacchi, M., Sharifi, M., Zeghidour, N., & Frank, C. (2023). MusicLM: Generating music from text. *arXiv preprint arXiv:2301.11325*.
18. Copet, J., Kreuk, F., Gat, I., Remez, T., Kant, D., Synnaeve, G., Adi, Y., & Défossez, A. (2024). Simple and controllable music generation. *Advances in Neural Information Processing Systems*, 36.
19. Borsos, Z., Marinier, R., Vincent, D., Kharitonov, E., Pietquin, O., Sharifi, M., Roblek, D., Teboul, O., Grangier, D., Tagliasacchi, M., & Zeghidour, N. (2023). AudioLM: A language modeling approach to audio generation. *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 31, 2523-2533.
20. Christiano, P. F., Leike, J., Brown, T. B., Martic, M., Legg, S., & Amodei, D. (2017). Deep reinforcement learning from human preferences. *Advances in Neural Information Processing Systems*, 30, 4299-4307.
21. Mittelstadt, B. D., Allo, P., Taddeo, M., Wachter, S., & Floridi, L. (2016). The ethics of algorithms: Mapping the debate. *Big Data & Society*, 3(2), 1-21.
22. Selbst, A. D., Boyd, D., Friedler, S. A., Venkatasubramanian, S., & Vertesi, J. (2019). Fairness and abstraction in sociotechnical systems. *Proceedings of the Conference on Fairness, Accountability, and Transparency*, 59-68.
23. Bender, E. M., Gebru, T., McMillan-Major, A., & Shmitchell, S. (2021). On the dangers of stochastic parrots: Can language models be too big? *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency*, 610-623.