



# Neural Methods for Data-to-text Generation

MANDAR SHARMA, AJAY KUMAR GOGINENI, and NAREN RAMAKRISHNAN, Virginia Tech, Blacksburg, VA, USA

---

The neural boom that has sparked natural language processing (NLP) research throughout the last decade has similarly led to significant innovations in data-to-text (D2T) generation. This survey offers a consolidated view into the neural D2T paradigm with a structured examination of the approaches, benchmark datasets, and evaluation protocols. This survey draws boundaries separating D2T from the rest of the natural language generation (NLG) landscape, encompassing an up-to-date synthesis of the literature, and highlighting the stages of technological adoption from within and outside the greater NLG umbrella. With this holistic view, we highlight promising avenues for D2T research that focus not only on the design of linguistically capable systems but also on systems that exhibit fairness and accountability.

CCS Concepts: • **Computing methodologies** → **Natural language generation**; *Learning paradigms*; *Machine learning algorithms*;

Additional Key Words and Phrases: Narration, data-to-text, data-to-text generation, natural language generation

## ACM Reference format:

Mandar Sharma, Ajay Kumar Gogineni, and Naren Ramakrishnan. 2024. Neural Methods for Data-to-text Generation. *ACM Trans. Intell. Syst. Technol.* 15, 5, Article 89 (October 2024), 46 pages.  
<https://doi.org/10.1145/3660639>

---

## 1 Introduction

*Textual representations of information: A picture is worth a thousand words—isn't it? And hence graphical representation is by its nature universally superior to text—isn't it? Why then isn't the anecdote itself represented graphically?*—Petre [246], in his advocacy for textual representation of information, challenges the notion that graphical representations of information are inherently more memorable, comprehensible, and accessible than their textual counterparts. Gershon and Page [104] note that the transformation of information from a textual to visual domain, in certain instances, requires further addition of information rendering textual representations more economical. Similarly, *knowing where to look* may not be obvious in visual representations of information—as validated through reading comprehension experiments [31, 115, 116] where participants were significantly slower in interpreting visual representations of the nested conditional structures within a program compared to their textual representations. This being said, these studies do not intend to dissuade the use of visual representations but rather establish the importance of textual representation of

---

This work is supported in part by US National Science Foundation grant DGE1545362.

Authors' Contact Information: Mandar Sharma (Corresponding author), Virginia Tech, Blacksburg, VA, USA; e-mail: mandarsharma@vt.edu; Ajay Kumar Gogineni, Virginia Tech, Blacksburg, VA, USA; e-mail: gogineniak20@vt.edu; Naren Ramakrishnan, Virginia Tech, Blacksburg, VA, USA; e-mail: naren@cs.vt.edu.




This work is licensed under a Creative Commons Attribution International 4.0 License.


© 2024 Copyright held by the owner/author(s).

ACM 2157-6912/2024/10-ART89

<https://doi.org/10.1145/3660639>



**T<sup>3</sup> Narrative (United Kingdom):** Coronavirus (initial case observed in January 2020) prevalence per year was reported to be 3515796 in the United Kingdom by January 2021. The virus invaded the UK in 2 waves. The first outbreak, peaked at 7156 cases, lasting from September 2020 to October 2020. The reported peak for the second outbreak that spanned October 2020 to January 2021 was 68192. The cases from January 2020 to May 2020 decreased to 502 while that from July 2020 to November 2020 rose from 12164 to 23000. In the time period from November 2020 to January 2021, cases increased to 68192. The number of cases decreased in January 2021 to 38992. The highest recorded Coronavirus daily case count in the UK was 68192 in January 2021.



**T<sup>3</sup> Narrative (Kansas):** The first trace of Carbon Monoxide as a measurement of the air quality was observed in January 2000 in the state of Kansas. The mean CO observed can be dissected into 2 regimes of interest. The first regime spanning April 2006 to May 2006 and had an average CO emission of 0.44. The second regime from April 2013 to April 2013 had an average of 0.13. From January 2000 to July 2002, the observed values decreased from 0.9 to 0.0. The observed values also decreased from 0.48 to 0.12 during February 2003 to May 2010. The mean CO emissions peaked in October 2000 at 2.58.



Fig. 1. Illustration of D2T: Narration of time series data (COVID19 progression in the United Kingdom at the top, the carbon monoxide emissions in the state of Kansas, United States, at the bottom) with the large language model (LLM)-based framework  $T^3$  (T-Cube) [291]. This D2T framework consumes a time series as input and generates narratives that highlight the progression and points-of-interest (regimes, trends, and peaks) in the data through LLM-generated narratives.

information. Often, the interplay of these paradigms brings out the best of both [289]. Thus, having established the importance of textual representations of information, we next explore how these notions tie into **data-to-text (D2T)** generation.

### 1.1 Defining D2T Generation and Scope of the Survey

Textual representations of information, for easier assimilation, are often presented as annotations outlining different behaviours of the underlying data stitched together. These stitched annotations, as showcased in Figure 1, are referred to as *narratives*. The automated generation of such narratives, although serving several niches (see below), are most prevalent in the public eye through the practice of robo-journalism [66, 176]. Bloomberg News generates a third of its content with Cyborg, their in-house automation system that can dissect tedious financial reports and churn out news articles within seconds [243]. Also prevalent are the use of such systems in business intelligence settings with prominent commercial frameworks,<sup>1</sup> such as Arria NLG,<sup>2</sup> Narrative Science,<sup>3</sup> and Automated Insights.<sup>4</sup>

The practice of automating the translation of data to user-consumable narratives through such systems is known as *D2T generation*, as depicted in Figure 1. Although encompassed by the general umbrella of *natural language generation (NLG)*, the nuance that differentiates D2T from the rest of the NLG landscape is that the input to the system has to qualify as a *data* instance. Reiter and Dale (1997) [270] describe the instance as a non-linguistic representation of information, and although narration of images and videos [70] has garnered interest in the NLG community, the definition of D2T employed by this survey follows that established by the seminal works prior [100, 270]: an

<sup>1</sup>This survey exclusively focuses on academic innovations for D2T generation as the technologies underlying commercial frameworks are often proprietary.

<sup>2</sup><https://www.arria.com/>

<sup>3</sup><https://narrativescience.com/>

<sup>4</sup><https://automatedinsights.com/>

entity that is *not exclusively linguistic*—tabular databases, graphs and knowledge bases, time series, and charts. Using this clause, we limit the scope of our analysis and exclude examination of all other NLG systems that either both ingest and expel linguistic entities for downstream tasks, such as **machine translation (MT)** [145, 350] and summarization [188, 228], or ingest non-conventional data, such as images [131] and videos [328].

Outside of dataset specific tasks, practical applications of D2T include, but are not limited to:

- Weather forecasts [20, 271]
- Sport summaries [15, 279, 316]
- Healthcare [241, 249]
- Virtual dietitians [9]
- Stock market comments [10, 220]
- Video game dialogues [149] and driving feedback [35]

With our scope defined, below we outline the rationale for this survey, followed by a structured examination of approaches, benchmark datasets, and evaluation protocols that constitute the D2T landscape with the intent to outline promising avenues for further research.

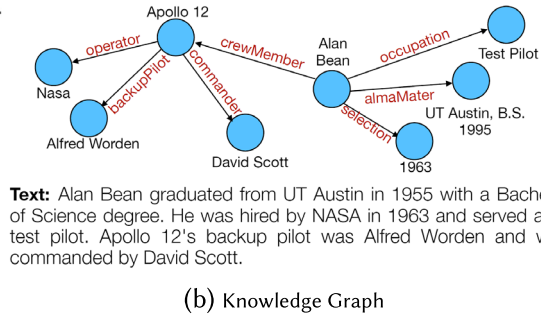

## 1.2 Survey Rationale

Following the seminal work by Reiter and Dale [270], the most comprehensive survey on D2T to date has been that by Gatt and Krahmer [100]. Although several articles have taken a close examination of NLG sub-fields such as dialogue systems [282], poetry generation [234], persuasive text generation [76], social robotics [94], or exclusively focus on issues central to NLG, such as faithfulness [181] and hallucination [142], a detailed breakdown of the last half-decade of innovations has been missing since the last exhaustive body of work. The need for a close and consolidated examination of developments in neural D2T is more pertinent now than ever. Further, D2T distinguishes itself from other NLG tasks as it blends the generation of narratives with numerical reasoning between data points. Outside of the D2T niche, there are research communities focused on solving these individual problems—NLG [100, 235, 320] and numerical reasoning [295, 317, 334, 349]. Thus, neural D2T is uniquely positioned such that it has either to incorporate innovations from these seemingly disparate niches or to jointly innovate on both fronts. We believe this provides added justification for D2T requiring its own comprehensive literature review.

As such, neural D2T borrows heavily from advances in other facets of NLG, such as **neural MT (NMT)** [11, 350] and **spoken dialogue systems (SDS)** [79, 342, 343]. As such, the pertinence of such a survey also spans highlighting the stages of technological adoptions in the D2T paradigm and drawing distinctions between its NMT and SDS neighbors. Further, the adoption of such technologies brings about the adoption of shared pitfalls—inconsistencies in evaluation metrics [268] and meaningful inter-model comparisons [265]. Thus, in addition to an exhaustive examination of neural D2T frameworks, a consolidated resource on approaches to its evaluation is also necessary. Also crucial, is the discussion of benchmark datasets across shared tasks. The above considerations motivate our survey on the neural D2T paradigm intended to serve the following goals:

- Structured examination of innovations in neural D2T in the last half-decade spanning relevant frameworks, datasets, and evaluation measures.
- Outlining the technological adoptions in D2T from within and outside of the greater NLG umbrella with the distinctions and shared pitfalls that lie therein.
- Highlighting promising avenues for further D2T research and exploration that promote fairness and accountability along with linguistic prowess.

Obama was elected and his voters celebrated



**Text:** Alan Bean graduated from UT Austin in 1955 with a Bachelor of Science degree. He was hired by NASA in 1963 and served as a test pilot. Apollo 12's backup pilot was Alfred Worden and was commanded by David Scott.

Fig. 2. AMR, Abstract MR [169] and knowledge graph [274] snapshots, representing variants of graph-based inputs to D2T systems.

## 2 Datasets for D2T Generation

The first set of technological adoptions from **natural language processing (NLP)** takes the form of *dataset design*: parallel corpora that align the data to their respective narratives are crucial for end-to-end learning, analogous to any neural-based approach to text processing. The initial push toward building such datasets began with database–text pairs of weather forecasts [20, 271] and sport summaries [15]. These datasets, and the convention that currently follows, use *semi-structured* data that deviates from the raw numeric signals initially used for D2T systems [266]. The statistics for prominent datasets among the ones discussed below are detailed in Table 2.

### 2.1 Meaning Representations

Mooney [217] defines a *meaning representation (MR) language* as a formal unambiguous language that allows for automated inference and processing wherein natural language is mapped to its respective MR through semantic parsing [101]. Robocup [42], among pioneering MR-to-text datasets, offers data from 1,539 pairs of temporally ordered simulated soccer games in the form of MRs (pass, kick, turnover) accompanied with their respective human commendation. To mitigate the cost of building large-scale MR datasets, Liang et al. [183] use grounded language acquisition to construct WeatherGov—a weather forecasting dataset with 29,528 MR–text pairs, each consisting of 36 different weather states. **Abstract MR (AMR)** [13], similarly, is a linguistically grounded semantic formalism representing the meaning of a sentence as a directed graph, as depicted in Figure 2(a). The LDC repository<sup>5</sup> hosts various AMR-based corpora. Following this, using simulated dialogues between their statistical dialogue manager [357] and an agenda-based user simulator [283], Mairesse et al. [204] offer BAGEL—an MR–text collection of 202 Cambridge-based restaurant descriptions each accompanied with two *inform* and *reject* dialogue types. Wen et al. [343], through crowdsourcing, offer an enriched dataset consisting of 5,192 instances of 6 additional dialogue act types, such as *confirm* and *inform only* (8 total) for hotels and restaurants in San Francisco. Novikova et al. [232] show that crowdsourcing with the aid of pictorial stimuli yield better phrased references compared to textual MRs. Following this, they released the E2E dataset<sup>6</sup> as a part of the E2E challenge [230]. With 50,602 instances of MR–text pairs of restaurant descriptions, its lexical richness and syntactic complexity provides new challenges for D2T systems. Table 1 showcases comparative snapshots of the aforementioned datasets.

<sup>5</sup><https://amr.isi.edu/>

<sup>6</sup><http://www.macs.hw.ac.uk/InteractionLab/E2E/>

Table 1. Comparative Showcase of Sample MRs (and Their Corresponding Narratives) from the RoboCup, WeatherGov, BAGEL, SF Hotels and Restaurants, and E2E Datasets

MR	Text
<i>RoboCup</i> [42] badPass(arg1=pink11,...), ballstopped() ballstopped(), kick(arg1=pink11) turnover(arg1=pink11,...)	pink11 makes a bad pass and was picked off by purple3
<i>WeatherGov</i> [183] rainChance(time=26-30,...), temperature(time=17-30,...) windDir(time=17-30,...), windSpeed(time=17-30,...) precipPotential(time=17-30,...), rainChance(time=17-30,...)	Occasional rain after 3am. Low around 43. South wind between 11 and 14 mph. Chance of precipitation is 80%. New rainfall amounts between a quarter and half of an inch possible.
<i>BAGEL</i> [204] inform( name(the Fountain) near(the Arts Picture House) area(centre), pricerange(cheap))	There is an inexpensive restaurant called the Fountain in the centre of town near the Arts Picture House
<i>SF Hotels and Restaurants</i> [343] inform( name="red door cafe", goodformeal="breakfast", area="cathedral hill", kidsallowed="no")	red door cafe is a good restaurant for breakfast in the area of cathedral hill and does not allow children.
<i>E2E</i> [232] name[Loch Fyne], eatType[restaurant], food[French], priceRange[less than £20], familyFriendly[yes]	Loch Fyne is a family-friendly restaurant providing wine and cheese at a low cost.

## 2.2 Graph Representations

Graph-to-text translation is not only central to D2T as its application carries over to numerous NLG fields such as question answering [74, 130], summarization [86], and dialogue generation [190, 216]. Further, the D2T frameworks for graph-to-text borrow heavily from the theoretic formulations offered from the literature in the field of graph neural networks, as will be discussed in Section 5.1.5. The *domain-specific* benchmark datasets, as discussed above (see Section 2.1) inherently train models to generate stereotypical domain-specific text. By crowdsourcing annotations for DBpedia [211] graphs spanning 15 domains, Gardent et al. [99] introduce the WebNLG dataset.<sup>7</sup> The data instances are encoded as *resource description format* triples of the form (*subject, property, and object*) as depicted in Figure 2(b)—(*Apollo 12, operator, NASA*). With 27,731 multi-domain graph-text pairs, WebNLG offers more semantic and linguistic diversity than previous datasets twice its size [342]. The abstract generation dataset [168], built with knowledge graphs extracted from articles in the proceedings of AI conferences [6] using SciIE [199], offers 40,000 graph-text pairs of the article abstracts. To further promote generation challenges and cross-domain generalization, Nan et al. [224] merge the E2E and WebNLG dataset with large heterogeneous collections of diverse predicates from Wikipedia tables annotated with tree ontologies to generate the data-record-to-text corpus. With 82,191 samples, this resulting open-domain corpus is almost quadruple the size of WebNLG.

## 2.3 Tabular Representations

Information represented in large tables can be difficult to comprehend at a glance, thus, **table-to-text (T2T)** aims to generate narratives highlighting crucial elements of a tabular data instance through summarization and logical inference over the table—as showcased in Figure 3. Similar to graph-to-text, the underpinnings of tabular representation learning is also shared with other fields outside of NLG, such as the generation of synthetic network traffic [276, 353].

WikiBio [174], as an initial foray toward a large-scale T2T dataset, offers 700k table-text pairs of Wikipedia info-boxes with the first paragraph of its associated article as the narrative. With a vocabulary of 400k tokens and 700k instances, WikiBio offers a substantially larger benchmark

<sup>7</sup><https://webnlg-challenge.loria.fr/>

TEAM	WIN	LOSS	PTS	FG.PCT	RB	AS ...
Heat	11	12	103	49	47	27
Hawks	7	15	95	43	33	20

PLAYER	AS	RB	PT	FG	FGA	CITY ...
Tyler Johnson	5	2	27	8	16	Miami
Dwight Howard	4	17	23	9	11	Atlanta
Paul Millsap	2	9	21	8	12	Atlanta
Goran Dragic	4	2	21	8	17	Miami
Wayne Ellington	2	3	19	7	15	Miami
Dennis Schroder	7	4	17	8	15	Atlanta
Rodney McGruder	5	5	11	3	8	Miami
Thabo Sefolosha	5	5	10	5	11	Atlanta
Kyle Korver	5	3	9	3	9	Atlanta
...						

The Atlanta Hawks defeated the Miami Heat , 103 - 95 , at Philips Arena on Wednesday . Atlanta was in desperate need of a win and they were able to take care of a shorthanded Miami team here . Defense was key for the Hawks , as they held the Heat to 42 percent shooting and forced them to commit 16 turnovers . Atlanta also dominated in the paint , winning the rebounding battle , 47 - 34 , and outscoring them in the paint 58 - 26. The Hawks shot 49 percent from the field and assisted on 27 of their 43 made baskets . This was a near wire - to - wire win for the Hawks , as Miami held just one lead in the first five minutes . Miami ( 7 - 15 ) are as beat - up as anyone right now and it 's taking a toll on the heavily used starters . Hassan Whiteside really struggled in this game , as he amassed eight points , 12 rebounds and one blocks on 4 - of - 12 shooting ...

Fig. 3. Showcasing the intent of T2T, the statistics of a basketball match between the Atlanta Hawks and the Miami Heat (left) is to be translated into easily consumable narratives (right). Snapshot from the RotoWire dataset [347].

compared to the pioneering WeatherGov and Robocup datasets that have less than 30k data-text pairs. For neural systems, as the length of output sequence increases, the generated summary diverges from the reference. As such, the RotoWire dataset [347] (Figure 3), consisting of verbose descriptions of NBA game statistics, brings forth new challenges in long-form narrative generation as the average reference length of RotoWire is 337 words compared to 28.7 of WikiBio. Similarly, with the observation that only 60% of the content in RotoWire narratives can be traced back to the data records, Wang [335] introduce RotoWire-FG, a refined version of the original dataset aimed at tackling divergence (see Section 3.2), where narrative instances not grounded by their respective tables are removed from the dataset.

TabFact [50] contains annotated sentences that are either supported or refuted by the tables extracted from Wikipedia. Similar to RotoWire-FG, Chen et al. [47] offer a filtered version of TabFact by retaining only those narratives that can be logically inferred from the table.

For controlled generation, Parikh et al. [239] propose ToTTo which generates a single sentence description of a table on the basis of a set of highlighted cells where annotators ensure that the target summary only contains the specified subset of information. With over 120k training samples, ToTTo establishes an open-domain challenge for D2T in controlled settings. Similarly, to evaluate narrative generation in open-domain settings with sentences that can be logically inferred from mathematical operations over the input table, Chen et al. [47] modify the reference narratives of the TabFact dataset to construct LogicNLG with 7392 tables. Following this, with tables and their corresponding descriptions extracted from scientific articles, Moosavi et al. [218] introduce SciGen, where the narratives include arithmetic reasoning over the tabular numeric entries. Building upon the long form generation premise of RotoWire, Chen et al. [45] construct WikiTableT, a multi-domain table-text dataset with 1.5 million instances pairing Wikipedia descriptions to their corresponding info-boxes along with additional hyperlinks, named-entities, and article metadata. The majority of these datasets are available in a unified framework through TabGenie.<sup>8</sup>

<sup>8</sup><https://pypi.org/project/tabgenie/>

Table 2. Highlights from Prominent D2T Datasets: Format, Number of Samples (Size), the Number of Linguistic Tokens Across the Dataset (Tokens), and Availability of Non-Anglo-Centric Variants

Benchmark	Format	Size	Tokens
E2E	MR	50,602	65,710
LDC2017T10	AMR	39,260	-
WebNLG (en, ru)	RDF	27,731	8,886
Data-record-to-text	RDF	82,191	33,200
WikiBio	Record	728,321	400,000
RotoWire	Record	4,853	11,300
TabFact	Record	16,573	-
ToTTo	Record	120,000	136,777
LogicNLG	Record	37,000	52,700
WikiTableT	Record	1.5M	169M

## 2.4 Data Collection and Enrichment

The majority of the prominent datasets discussed in Sections 2.1–2.3 are either collected by merging aligned data–narrative pairs that occur naturally in the “wild” [174, 347] or collected through dedicated crowd-sourcing approaches [99, 230]. However, there are notable works that employ a hybrid approach to data collection. CACAPO [324], an MR-style multi-domain dataset, follows a collection process inspired by Oraby et al. [236] wherein the naturally occurring narratives are first scraped from the internet and are later manually annotated to generate attribute–value pairs. Similarly, chart-to-text [153] follows a similar mechanism of data collection wherein candidate narratives for each chart are first automatically generated via a heuristic-based approach and then are rated by crowd-sourced workers. In similar lines, the ToTTo dataset [239] discussed in Section 2.3 uses crowd-sourced annotators as data “cleaners”—iteratively improving upon the automatically scraped narratives, rather than annotating them from scratch—thus greatly reducing the cost of data acquisition. In addition to innovations in data collection, efforts from the D2T community have also focused on the enrichment of existing datasets. As such, Ferreira et al. [90] augment the WebNLG dataset with intermediate representation for discourse ordering and referring expression generation. By manually delexicalizing (see Section 4.1) the narratives, Ferreira et al. were able to automatically extract a collection of referring expressions by tokenizing the original and delexicalized texts and finding the non-overlapping tokens between them. Similarly, the authors also extracted the order of the arguments in the text by referring to the order of the general tags in the delexicalized texts. This work has also been extended to enrich the E2E dataset [92].

## 3 D2T Generation Fundamentals and Notations

### 3.1 What to Say and How to Say It

The data instance, typically, contains more information than what we would intend for the resulting narrative to convey—verbose narratives that detail every attribute of the data instance contradicts the premise of consolidation. Thus, to figure out *what to say*, a subset of the original information content is filtered out based on the target audience through the process of **content selection (CS)**. Starting from data-driven approaches such as clustering [75] and the use of hidden Markov models [16], the attention of the research community has recently shifted to learning *alignments* between

the data instance and its narrative [183]. Bisazza and Marcello [30] note that pre-reordering the source words to better resemble the target narrative yields significant improvements in NMT. Prior to neural explorations, learning this alignment has been explored with log-linear models [8] and tree representations [170, 171]. With *what to say* determined, the next step lies in figuring out *how to say it*, that is, the construction of words, phrases, and paragraphs—this realization of the narrative structure is known as *surface realization*. While traditionally, the processes of CS and surface realization [148, 270] act as discrete parts of the generation pipeline, the neural **sequence-to-sequence (seq2seq)** paradigm jointly learns these aspects. For a peripheral view of the articles discussed in this section, Table 4 highlights prominent papers categorized based on their D2T tasks and the benchmark datasets used. Similarly, Figure 5 outlines the organization of the remainder of this survey.

### 3.2 Hallucinations and Omissions

Apart from the importance of coherence and linguistic diversity in surface realization, *data fidelity* is a crucial aspect of D2T systems—the narrative should neither *hallucinate* contents absent from the data instance nor *omit* contents present in the data instance. Often, the divergence present in benchmark training datasets, wherein the narrative may contain data absent from the source or not cover the entirety of the data instance, is the culprit behind hallucination tendencies in the model [280]. Often times, the need for both linguistic diversity and data fidelity turns into a balancing act between conflicting optimization objectives leading to novel challenges [128]. While almost all of the D2T approaches discussed below engage in balancing coherence and diversity with data-fidelity (besides Section 5.1.4), overarchingly, the approach to balancing these conflicting objectives can be thought to take place in two forms:

- *Architectural interventions*: The Sections 5.1.1–5.1.3, 5.1.5, 5.1.6, and 5.1.10 suggest modifications or augmentations to the seq2seq architecture such that it fosters data-fidelity tendencies.
- *Loss-function interventions*: An alternative avenue to achieving a balance between conflicting optimization objectives is to directly model the objective functions to perform multi-task learning: as such Sections 5.1.7 and 5.1.8 suggest modifications or augmentations to the seq2seq loss functions.

### 3.3 Establishing Notation and Revisiting Seq2Seq

For the consistency and readability of this survey, the notation outlining the basic encoder–decoder seq2seq paradigm [11, 54, 314, 326] in D2T (Figure 4), as defined below and respectively compiled in Table 3, will remain valid throughout unless stated otherwise. However, the namespace for additional variable definitions in the individual sections will be limited to their mentions. Let  $S = \{x_j, y_j\}_{j=1}^N$  be a dataset of  $N$  data instances  $x$  accompanied with its natural language narrative  $y$ . Based on the construction of  $S$ ,  $x$  can be a set of  $K$  data records  $x = \{r_j\}_{j=1}^K$  with each entry  $r$  comprised of its respective entity  $r.e$  and value  $r.m$  attributes or  $x$  can be an instance of a directed graph  $x = (V, E)$  with vertices  $v \in V$  and edges  $(u, v) \in E$ . In the RotoWire instance (Figure 3), for  $r_j = \text{Heat}$ ,  $r_j.e = \text{WIN}$  attribute would have value  $r_j.m = 11$ . Given pairs  $(x, y)$ , the seq2seq model  $f_\theta$  is trained end-to-end to maximize the conditional probability of generation  $P(y|x) = \prod_{t=1}^T P(y_t | y_{<t}, x)$ . The parameterization of  $f_\theta$  is usually carried out through RNNs, such as LSTMs [29, 133] and GRUs [54], or transformer<sup>9</sup> architectures [326]. For attention-based RNN

<sup>9</sup>Though the base transformer architecture is oblivious to input structures, we assume positionally encoded transformers to fall into the seq2seq paradigm.



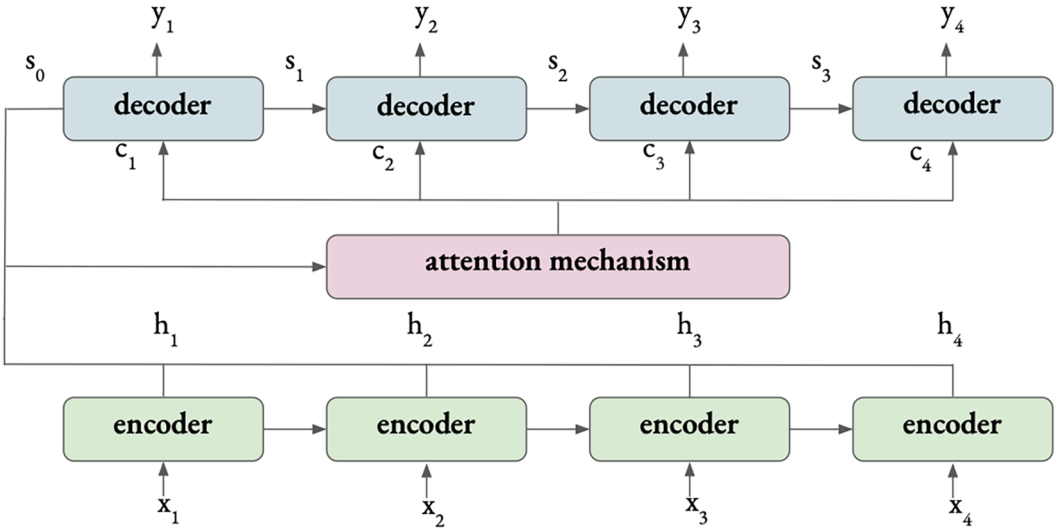


Fig. 4. Attention-based seq2seq framework: The encoder consumes the sequential input translating it to a weighed hidden representation to be then consumed and decoded into linguistic tokens by the decoder.

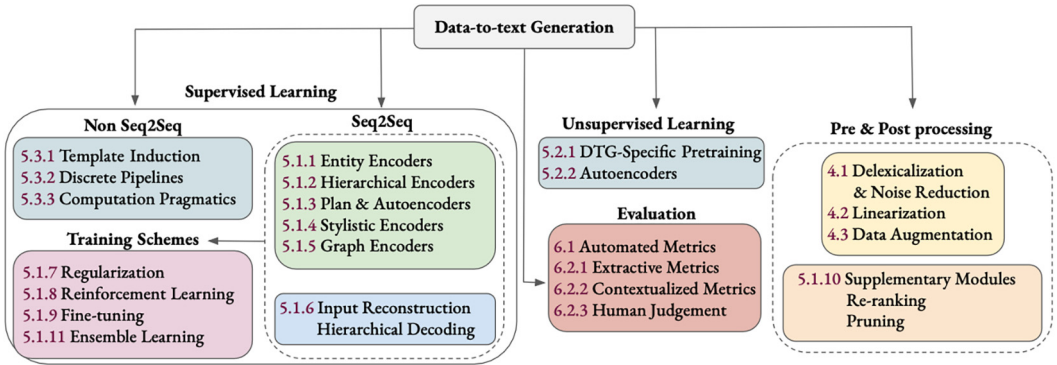


Fig. 5. D2T generation taxonomy corresponding to sections in the survey design.

architectures with hidden states  $h_t$  and  $s_t$  for the encoder and decoder respectively, the context vector  $c_t = \sum_i \alpha_{t,i} h_i$  weighs the encoder hidden states with attention weights  $\alpha_{t,i}$ . While Bahdanau et al. [11] use a **multi-layer perceptron (MLP)** to model  $\alpha_{t,i}$ , several alterations to modeling the attention weights have been proposed [201, 331].

For handling of out-of-vocabulary tokens, Gu et al. [119] attempt to model the rote memorization process of human learning where a language model conditioned on binary variable  $z_t \in \{0, 1\}$  can either generate  $p_{gen}$  the next token or copy it from the source  $p_{copy}$  based on their respective probabilities. While Gu et al. [119] and Yang et al. [355] parameterize the *joint* distribution over  $y_t$  and  $z_t$  directly (Equation (1)), Gulçehre et al. [121] decompose the joint probability (Equation (2)),

Table 3. Notation Descriptions

Notation	Description
$S$	Dataset
$(x, y) \in S$	Data instance $x$ and its natural language representation $y$
$r = (r.e, r.m)$	Data record $r$ with its entity $r.e$ and value $r.m$ attributes
$G = (V, E)$	Graph instance $G$ with vertices $V$ and edges $E$
$(u, v) \in E$	Nodes $u$ and $v$ of an edge $E$
$f_\theta \in \{f_1, \dots, f_n\}$	Model $f_\theta$ that may belong to an ensemble $\{f_1, \dots, f_n\}$
$P(y x)$	Conditional probability of sequence $y$ given $x$
$h_t, s_t$	Encoder $h_t$ and decoder $s_t$ hidden state representations
$c_t, \alpha_{t,i}$	Context vector $c_t$ weighing $h_t$ with attention weights $\alpha_{t,i}$
$p_{gen}, p_{copy}$	Token generation $p_{gen}$ or copying $p_{copy}$ probabilities
$z_t \in \{0, 1\}$	Binary variable that selects either $p_{gen}$ or $p_{copy}$
$W_{i \in \mathbb{N}}, b_{i \in \mathbb{N}}$	Arbitrary weights and biases parameterizing $f_\theta$

using an MLP to model  $p(z_t|y_{<t}, x)$

$$P(y_t, z_t|y_{<t}, x) \propto \begin{cases} p_{gen}(y_t, y_{<t}, x)z_t = 0 \\ p_{copy}(y_t, y_{<t}, x)z_t = 1, y_t \in x \\ 0 & z_t = 1, y_t \notin x \end{cases} \quad (1)$$

$$\begin{cases} p_{gen}(y_t|z_t, y_{<t}, x)p(z_t|y_{<t}, x)z_t = 0 \\ p_{copy}(y_t|z_t, y_{<t}, x)p(z_t|y_{<t}, x)z_t = 1 \end{cases} \quad (2)$$

Similar to the greater NLG paradigm, different strategies for modeling the conditional probability of generation  $P(y|x)$ , the attention mechanisms  $\{\alpha_{t,i}, c_t\}$ , and the copy mechanisms  $\{p_{gen}, p_{copy}\}$ , as discussed below, *often form the basis for D2T innovations*. In addition to this, variations in training strategies such as teacher-forcing [346], **reinforcement learning (RL)** [315], and auto-encoder-based reconstruction [41] open up further avenues for D2T innovation.

## 4 Innovations in Data Preprocessing

Contrary to the other facets of NLG, such as chatbots, for which large-scale data can be harvested [1, 198], D2T datasets are often smaller in scale and task-specific. Ferreira et al. [88] note that phrase-based translation models [166] can outperform neural models in such data sparsity. As such, *delexicalization, noise reduction, linearization, and data augmentation* are preprocessing techniques often employed to tackle said sparsity of training data.

### 4.1 Delexicalization and Noise Reduction

*Delexicalization*, often referred to as *anonymization*, is a common practice in D2T [79, 204] wherein the slot-value pairs for the entities and their attributes in training utterances are replaced with a placeholder token such that weights between similar utterances can be shared [227]—as illustrated in Figure 6(a). These placeholder tokens are later replaced with tokens copied from the input data instance [174]. In comparison to copy-based methods for handling rare entities, delexicalization has shown to yield better results in constrained datasets [300].

From the notion that delexicalization of the data instance may cause the loss of vital information that can aid seq2seq models in sentence planning, where some data instance slots may even be deemed nondelexicalizable [343], Nayak et al. [227] explore different nondelexicalized




Fig. 6. Illustrations of delexicalization in MRs [300] and linearization of graphs [274].

input representations (mention representations) along with grouping representations as a form of sentence planning (plan representations). The authors note improvements over delexicalized seq2seq baselines when input mentions are concatenated with each slot–value pair representing a unique embedding. The efficacy of such concatenation is also corroborated by Freitag and Roy [95]. Further, the addition of positional tokens representing intended sentence position to the input sequence offers further improvements. Addressing this, in addition to delexicalizing categorical slots, Juraska et al. [150] employ hand-crafted tokens for values that require different treatment in their verbalization: for the slot *food*, the value *Italian* is replaced by *slot\_vow\_cuisine\_food* indicating that the respective utterance should start with a vowel and the value represents a cuisine—an *Italian restaurant*. Perez-Beltrachini and Lapata [244] delexicalize numerical expressions, such as dates, using tokens created with the attribute name and position of the delexicalized token. Colin and Gardent [59] note performance improvements with an extensive anonymization scheme wherein all lemmatized content words (expect adverbs) are delexicalized as compared to restricting delexicalization to named entities.

The presence of narratives that fail to convey vital attributes of data instances leads to *semantic noise* in the dataset [84, 136]. Dušek et al. [78] employ slot matching [263] to clean the E2E corpus for semantic correctness and explore the impact of semantic noise on neural model performance. Similarly, Obeid and Hoque [233] substitute data mentions in the narrative, identified through **named entity recognition (NER)**, with a predefined set of tokens which are later replaced through a look-up operation. Liu et al. [194] focus on generating faithful narratives and truncate the reference narratives by retaining only the first few sentences, since the latter are prone to have been *inferred* from the table.

#### 4.2 Linearization

Frameworks for MR (and graph) narration that shy from dedicated graph encoders rely on effective *linearization* techniques—the representation of graphs as linear sequences, as illustrated in Figure 6(b). While Ferreira et al. [88] note improvements in neural models with the adoption of a two-step classifier [177] that maps AMRs to the target text, Konstas et al. [169] showcase agnosticism to linearization orders by grouping and anonymizing graph entities for delexicalization with the Stanford NER [93]. The reduction in graph complexity and subsequent mitigation of the challenge brought forth by data sparsity lends any depth-first traversal of the graph as an effective linearization approach. Moryossef et al. [219] append text plans modeled as ordered trees [308] to the WebNLG training set and use an off-the-shelf NMT system [121] for plan-to-text generation. However, the authors note that the restriction of requiring single entity mentions in a sentence establishes their approach as dataset dependent.

For pretrained language models, such as GPT-2 [257] and T5 [258], Zhang et al. [361] and Gong et al. [109] represent tables as a linear sequence of attribute–value pairs and use a special token as the separator between the table data and the reference text. It should be noted that T5, while performing the best on automated metrics, fails to generate good summaries when numerical calculations are involved. Chen et al. [47] traverse the table horizontally, each row at a time, where each element is

represented by its corresponding field and cell value separated by the keyword *is*. For scientific tables, Suadaa et al. [312] view a table  $T_D$  as a set of cells with their corresponding row and column headers  $h = [rh : ch]$  with  $th$  for overlapping tokens, numerical value  $val$ , and metric-type  $m$ . The cells are marked with target flag  $tgt$  which is set to 1 for *targeted cells* and 0 otherwise respective to the content plan. The linearization of the resulting tables is done with templates that consist of concatenation  $T_D = [h : th : val : m : tgt]$ , filtration based on  $tgt$ , pre-computed mathematical operations, and their respective combinations.

### 4.3 Data Augmentation

Often, appending *contextual* examples from outer sources to the training set, or *permuting* the training samples themselves to append variation, helps mitigate data sparsity. This is known as data *augmentation*. Nayak et al. [227] propose the creation of pseudo-samples by permuting the slot orderings of the MRs while keeping the utterances intact. Juraska et al. [150], however, take an utterance-oriented approach where pseudo-samples are built by breaking training MRs into single-sentence utterances. For the shared surface realization task [213], Elder and Hokamp [83] augment the training set with sentences from the WikiText corpus [212] parsed using UDPipe [309]. Following this premise, Kedzie and Mckeown [159] curate a collection of utterances from novel MRs using a vanilla seq2seq model with noise injection sampling [53]. The validity of the MRs associated with these utterances are computed through a CNN-based parser [162] and the valid entries are augmented to the training set. However, it is worth noting that performance gains from augmenting the training set with out-of-domain instances tend to saturate after a certain point [95]. Also, practitioners of data augmentation should note that caution is advised when augmenting with synthetic data, as the inclusion of such data may reinforce the mistakes of the model [126].

Chen et al. [46] append knowledge graphs representing *external context* to the table–text pairs and quantify its efficacy through their metric KBGain—the ratio of tokens unique to the external context to the total number of tokens in the narrative. Similarly, Ma et al. [202] augment the limited training data for table–text pairs by assigning **part-of-speech (POS)** tags for each word in the reference and further increase the robustness of their model with adversarial examples created by randomly adding and removing words from the input. In contrast, Chen et al. [47] create adversarial examples by randomly swapping entities in the narrative with ones that appear in the table. Following this, Liu et al. [194] use an augmented plan consisting of table records and entities recognized from the reference narrative which eliminates the inclusion of information not present in the table. For few-shot learning, Liu et al. [189] observed that the performance of a GPT-3 model [37] improved upon providing in-context examples computed based on their  $k$ -nearest neighbor ( $k = 2$ ) embeddings.

## 5 Innovations in the Seq2Seq Framework

Seq2Seq models (see Section 3.3) serve as the basis for neural NLG [54, 314, 326]. As such, to compare the efficacy of neural architectures for long-form D2T, Wiseman et al. [347] compare the performance of various seq2seq models to their templated counterparts on the RotoWire dataset. Based on their observations, the conditional copy model [121] performs the best on both word-overlap and *extractive metrics* (see Section 6.2.1) compared to the standard attention-based seq2seq model [11] and its joint copy variant [119]. Similarly, in an evaluation of 62 seq2seq, data-driven, and templated systems for the E2E shared task, Dušek et al. [81] note that seq2seq systems dominate in terms of both automated word-based metrics and *naturalness* in human judgment. Wiseman et al. [347], however, note that the traditional templated generation models outperform seq2seq models on extractive metrics although they score poorly on word-overlap metrics. Thus, the adaptation of

Table 4. Task and Dataset Based Summarization of Noted D2T Frameworks over the Last Half-Decade

Dataset	Publication Highlights	Framework and Human Evaluation
<b>MR-to-Text</b>		
Robocup and WeatherGov	[210] Coarse-to-fine aligner and penalty based on learned priors	LSTM → LSTM + regularization N
Recipe and SF H&R	[160] Neural agenda-checklist modeling	GRU → GRU + agenda encoders Y
BAGEL	[79] Reranking beam outputs w/ RNN-based reranker	LSTM → LSTM + reranker N
Restaurant Ratings	[227] Nondelexicalized inputs w/ data augmentation	LSTM → LSTM Y
WikiData	[52] Complementary text-to-data translation	GRU → GRU Y
E2E	[81] Comparative evaluation of 62 systems	Seq2Seq + Data-driven + Templated Y
	[256] MLP encoder attuned to the dataset	MLP → GRU Y
	[360] Two-level hierarchical encoder	CAEncoder → GRU N
	[150] Ensemble w/ heuristic reranking	Ensemble w/ LSTM + CNN N
	[310] Hierarchical decoding with POS tags	GRU → GRU N
	[95] Unsupervised DTG with DAEs	LSTM → LSTM + DAEs Y
	[103] Comparative evaluations w/ ensembling & penalties	Ensemble w/ LSTM + T N
	[62] Syntactic controls with SC-LSTM	SC-LSTM Y
	[297] Computational pragmatics based DTG	GRU → GRU N
	[59] Extensive anonymization	LSTM → LSTM Y
	[78] Semantic correctness in neural DTG	LSTM → LSTM Y
	[159] Self-training w/ noise injection sampling	GRU → GRU Y
	[277] Char-level GRU w/ input reconstruction	GRU → GRU N
[97] CRFs w/ Gumbel categorical sampling	CRF N	
<b>Graph-to-Text</b>		
AGENDA	[168] Graph-centric Transformer and AGENDA dataset	T → LSTM + LSTM encoding Y
LDC2015E25	[88] Phrase vs Neural MR-text w/ preprocessing analysis	LSTM → LSTM + Phrase-based N
LDC2015E86	[169] Unlabeled pre-training and linearization agnosticism	LSTM → LSTM N
	[273] Dual encoding for hybrid traversal	GNN → LSTM Y
LDC2017T10	[12] Graph reconstruction w/ node and edge projection	T → T + reconstruction loss Y
	[203] Fine-tuning GPT-2 on AMR-text joint distribution	GPT-2 Y
WebNLG	[207] Graph encoding with GCNs	GCN → LSTM N
	[68] LSTM based triple encoder	LSTM → LSTM Y
	[91] Discrete neural pipelines and comparisons to end-to-end	GRU → GRU + T Y
	[219] Sentence planning with ordered trees	LSTM → LSTM Y
	[275] Complementary graph contextualization	GAT → T Y
	[306] Detachable multi-view reconstruction	T → T N
	[364] Dual encoder for structure and planning	GCN → LSTM Y
	[274] Task-adaptive pretraining for PLMs	BART + T5 Y
	[2] Knowledge enhanced language models and KeLM dataset	T5 Y
[158] Graph-text joint representations and pretraining strategies	BART + T5 Y	
<b>Record-to-Text (Table-to-text)</b>		
WikiBio	[174] Tabular positional embeddings and WikiBio dataset	LSTM → LSTM + Kneser-Ney N
	[14] Encoding tabular attributes and WikiTableText dataset	GRU → GRU N
	[193] Field information through modified LSTM gating	LSTM → LSTM N
	[292] Link-based and content-based attention	LSTM → LSTM N
	[244] Multi-instance learning w/ alignment-based rewards	LSTM → LSTM Y
	[202] Key fact identification and data augument for few shot	LSTM + T N
	[191] Hierarchical encoding w/ supervised auxiliary learning	LSTM → LSTM Y
	[192] Forced attention for omission control	LSTM → LSTM Y
	[46] External contextual information w/ knowledge graphs	GRU → GRU Y
	[318] Confidence priors for hallucination control	BERT + Pointer Networks Y
	[51] Soft copy switching policy for few-shot learning	GPT-2 Y
	[356] Variational auto-encoders for template induction	VAE modified to VTM Y
	[337] Autoregressive modeling with iterative text-editing	Pointer networks + Text editing Y
	[365] Reinforcement learning with adversarial networks	GAN N
	[105] Linearly combined multi-reward policy	Pointer networks N
	[354] Source-target disagreement auxiliary loss	T N
	[311] BERT-based IR system for contextual examples	T5 + BERT Y
Rotowire	[347] Classification-based metrics and RotoWire dataset	LSTM → LSTM + Templated N
	[229] Numeric operations and operation-result encoding	GRU → GRU + operation encoders Y
	[253] Dynamic hierarchical entity-modeling and MLB dataset	LSTM → LSTM Y
	[252] Content selection and planning w/ gating and IE	LSTM → LSTM Y
	[107] Contextualized numeric representations	LSTM → LSTM Y
	[139] Dynamic salient record tracking w/ stylized generation	GRU → GRU N
	[261] Two-tier hierarchical input encoding	T → LSTM N
	[180] Auxiliary supervision w/ reasoning over entity graphs	LSTM + GAT Y
	[255] Paragraph-centric macro planning	LSTM → LSTM Y
[254] Interweaved plan and generation w/ variational models	LSTM → LSTM Y	

(Continued)

Table 4. (Continued)

Dataset	Publication Highlights	Framework and Human Evaluation
<b>Record-to-Text (Continued)</b>		
TabFact	[47] Coarse-to-fine two-stage generation	LSTM + T + GPT-2 + BERT Y
WikiPerson	[339] Disagreement loss w/ optimal transport matching loss	T Y
Humans, Books and Songs	[109] Attribute prediction-based reconstruction loss	GPT-2 Y
ToTTo	[189] Contextual examples through k nearest neighbors	GPT-3 Y
	[312] Targeted table cell representation	GPT-2 Y
	[49] Semantic confounders w/ Pearl's do-calculus	DCVED + GPT Y
	[187] Table-to-logic pretraining for logic text generation	T5 + BART Y
LogicNLG	[223] Faithful generation with unlikelihood and replacement detection	T5 Y
NumericNLG	[44] Table serialization and structural encoding	T → GPT-2 Y
	[7] T5 infused with tabular embeddings	T5 N
<b>Cross-domain</b>		
E2E	[140] Char-based vs word-based seq2seq	GRU → GRU Y
	[348] Template induction w/ neural HSMM decoder	HSMM N
	[98] Training w/ partially aligned dataset	T → T + supportiveness Y
WebNLG	[157] Iteratively editing templated text	GPT-2 + LaserTagger N
DART	[127] RoBERTa-based semantic fidelity classifier	GPT-2 + RoBERTa Y
WikiBio	[48] Knowledge-grounded pre-training and KGTEXT dataset	T → T + GAT N
RotoWire	[186] Hybrid attention-copy for stylistic imitation	LSTM + T Y
WITA	[351] Disambiguation and stitching with PLMs	GPT3 + T5 N
	[77] Unified learning of D2T and T2D	T5 + VAE N
	[144] Search and learn in a few-shot setting	T5 + Search and Learn Y
<b>Time series-to-text</b>		
WebNLG and DART	[294] Open-domain transfer learning for time-series narration	BART + T5 + Time series analysis Y
<b>Chart-to-text</b>		
Chart2Text	[233] Preprocessing w/ variable substitution	T → T Y
Chart-to-text	[153] Neural baselines for Chart-to-text dataset	LSTM + T + BART + T5 Y

seq2seq models to D2T for richer narratives with less omissions and hallucinations still remains an active focus of the research community.

It is worth noting that all seq2seq models discussed below operate at the *word* level. Models operating at the *character* level [3, 112, 277] have shown reasonable efficacy with the added computational savings from forgoing the preprocessing steps of delocalization and tokenization. However, the attention garnered by them from the research community is slim. From their comparative analysis, Jagfeld et al. [140] note that as character-based models perform better on the E2E dataset while word-based models perform better on the more linguistically challenging WebNLG dataset, it is hard to draw conclusions on the framework most suited for generic D2T. In the sections that follow, we detail notable innovations over the last half-decade in seq2seq modeling, branched on the basis of their training strategies—*supervised* and *unsupervised* learning.

## 5.1 Supervised Learning

### 5.1.1 Entity Encoders.

Centering theory [118], as well as many other noted linguistic frameworks [40, 57, 125, 156, 172, 251], highlights the critical importance of *entity mentions* to the coherence of the generated narrative. The ordering of these entities (*r.e* in Section 3.3) is crucial for such narratives to be considered as *entity coherent* [154]. Unlike typical language models which are conditioned solely based on previously generated tokens  $c_t$ , Lebre et al. [174] provide additional context  $\{z_{c_t}, g_f, g_w\}$  to the generation where  $z_{c_t}$  represents table entity  $c_t$  as a triplet of its corresponding field name, start, and end positions, and  $\{g_f, g_w\}$  are one-hot encoded vectors where each element indicates the presence of table entities from the fixed field and word vocabularies—illustrated in Figure 7. Similarly, Bao et al. [14] encode the table cell  $c$  and attributes  $a$  as the concatenation  $[e_i^c : e_i^a]$  where the decoder uses this vector to compute the attention weights. Liu et al. [193] modify the LSTM unit with a *field* gate to update the cell memory indicating the amount of entity field information to be retained in the cell memory. Following [174], Ma et al. [202] use a Bi-LSTM to encode the concatenation of word, attribute and position embeddings. However, to indicate whether

Table ( $\mathbf{g}_f, \mathbf{g}_w$ )		Input Text					
Name	John Doe		John	Doe	(	18	April
Birthdate	18 April 1352	$\mathbf{c}_t$	13944	UNK	17	37	92
		$\mathbf{z}_{et}$	(name,1,2)	(name,2,1)	$\phi$	(birth.,1,3)	(birth.,2,2)

Fig. 7. Entity encoding scheme—Lebret et al. [174].

an entity is a *key fact*, a MLP classifier is used on said representation for binary classification. Inspired from Liu and Lapata [195], Gong et al. [108] construct a historical timeline by sorting each table record with respect to its date field. Three encoders encode a table entity separately in row  $r_{i,j}^r$ , column  $r_{i,j}^c$ , and time  $r_{i,j}^t$  dimensions. The concatenation of these representations are fed to a MLP to obtain a general representation  $r_{i,j}^{gen}$  over which specialized attention weights are computed to obtain the final record representation as  $\hat{r}_{i,j} = \alpha_r r_{i,j}^r + \alpha_c r_{i,j}^c + \alpha_t r_{i,j}^t$ . Exploiting the attributes of the E2E dataset—the set number of unique MR attributes and the limited diversity in lexical instantiations of their values, Puzikov and Gurevych [256] employ a simple approach wherein the recurrent encoder is replaced with one dense layer that takes in MR representations through embedding lookup. Similarly, to keep track of entity mentions in the SF dataset for long-form text generation, Kiddon et al. [160] introduce a *checklist* vector  $a_t$  that aids two additional encoders to track used (mentioned in the resulting narrative) and new (not mentioned as of time step  $t$ ) items on the defined agenda. The output hidden state is modeled as a linear interpolation between the three encoder states— $c_t^{gru}$  of the base GRU and  $\{c_t^{new}, c_t^{used}\}$  from the agenda models, weighted by a probabilistic classifier. Extending this concept of entity encoding to transformer-based architectures, Chen et al. [44] adapt the multi-headed attention layer architecture [326] to encode serialized table attributes that is then fed to a GPT-based decoder. Similarly, as a unified text-to-text alternative approach to [44], Andrejczuk et al. [7] include the row and column embeddings  $\hat{r}_{i,j}$  of the input table on top of the token embeddings for table-structure learning in a T5 model.

Certain D2T tasks, such as sport commentaries [15, 279, 316], require reasoning over *numeric* entities present in the input data instance. Although numeracy in language modeling is a prominent niche of its own [295, 317, 334], notable D2T-specific approaches include that of Nie et al. [229]—precomputing the results of numeric operations  $op_i \in \{\text{minus}, \text{argmax}\}$  on the RotoWire dataset, the authors propose the combination of dedicated *operation* and *operation-result* encoders, the latter utilizing a quantization layer for mapping lexical choices to data values, in addition to a record encoder. In similar fashion to [355], the concatenated embeddings  $\{r.idx, r.e, r.m\}$  fed to a bi-directional GRU generate record representations while the concatenated embeddings of  $op_i$  attributes fed to a non-linear layer yields operation representations. To address the difficulty in establishing lexical choices on sparse numeric values [271, 303], the authors add quantization to the operation-results encoder that maps results of scalar operations  $e$  (minus) to  $l \in L$  possible bins through a weighed representation ( $h_i = \sum_l \mu_{i,l} e$ ) using softmax scores of each individual result  $\mu_{i,l}$ . Following this body of work, to contextualize numeric representations and thus understand their logical relationships, Gong et al. [107] feed raw numeric embeddings for all numerals corresponding to the same table attributes to a transformer-based encoder to obtain their *contextualized* representations. Through a ranking scheme based on a fully connected layer, these contextualized representations are further trained to favor larger numbers.

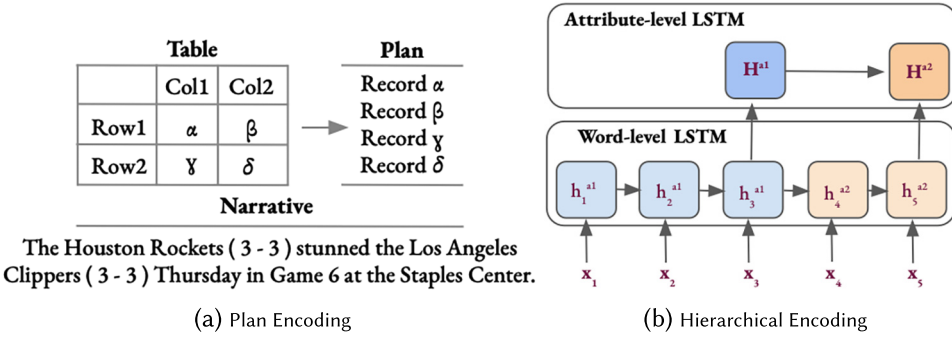


Fig. 8. Illustrations for plan [252] and hierarchical [191] encoding schemes.

**5.1.2 Hierarchical Encoders.** The intuition behind the use of hierarchical encoders, in the context of D2T, is to model input representations at different granularities, either through dedicated modules [191, 261, 360] or attention schemes [193, 253, 355]. As such, Zhang et al. [360] leverage their CAEncoder [359] to incorporate precomputed future representations  $h_{i+1}$  into current representation  $h_i$  through a two-level hierarchy. Similarly, Rebuffel et al. [261] propose a two-tier encoder to preserve the data structure hierarchy—the first tier encodes each entity  $e_i$  based on its associated record embeddings  $r_{i,j}$  while the second tier encodes the data structure based on its entity representation  $h_i$  obtained through the individual embeddings  $r_{i,j}$ . On the other hand, Liu et al. [191], as illustrated in Figure 8(b), propose a word-level  $h_{r,e}^{r,m}$  and an attribute-level  $H^{r,e}$  two-encoder setup to capture the attribute–value hierarchical structure in tables. The attribute-level encoder takes in the last hidden state  $h_{last}^{r,e}$  for attribute  $r.e$  from the word level LSTM as its input. Using these hierarchical representations, fine-grained attention  $\beta_{r,e}^{r,m} \propto g(h_{r,e}^{r,m}, s_t)$  and coarse-grained attention  $\gamma^{r,e} \propto g(H^{r,e}, s_t)$  are used for decoding where  $g$  represents a softmax function. Similarly, based on hierarchical attention [355], Liu et al. [193] employ an attention scheme that attends to both word level and field level tokens. Following this, Puduppully et al. [253] propose language modeling conditioned on both the data instance and a dynamically updated entity representation. At each time-step  $t$ , a gate  $\gamma_t$  is used to decide whether an update is necessary for the entity memory representation  $u_k$  and a parameter  $\delta_{t,k}$  decides the impact of said update (Equation (3))

$$\gamma_t = \sigma(W_1 s_t + b_1) \quad \& \quad \delta_{t,k} = \gamma_t \odot \sigma(W_2 s_t + W_3 u_{t-1,k} + b_3). \quad (3)$$

**5.1.3 Plan Encoders and Auto-Encoders.** Traditionally, the *what to say* aspect of D2T (see Section 3.1) used to be its own module in a set of pipelines [100, 270], thus offering flexibility in planning the narrative structure. However, the end-to-end learning paradigm often models CS and surface realization as a shared task [174, 210, 347]. Although convenient, without *explicitly* modeling the planning of the narrative (Figure 8(a)), language models struggle to keep coherence in long-form generation tasks. As such, Puduppully et al. [252] model the generation probability  $P(y|r)$  as the joint probability of narrative  $y$  and content plan  $z$  given a record  $r$  such that  $P(y|r) = \sum_z P(z|r)P(y|z)$ . Similar to their prior work [253], a CS gate operates over the record representation  $r_i$  giving an information controlled representation  $r_i^{cs}$ . The elements of  $z$  are extracted using an information extraction system [347] and correspond to entities in  $y$  while pointer networks [331] are used to align elements in  $z$  to  $r$  during training. Iso et al. [139], on the other hand, avoid precomputing content plans  $z$  by dynamically choosing data records during decoding—an additional memory state  $h^{ent}$  remembers mentioned entities and updates the language model state  $h^{lm}$  accordingly. The authors propose two representations for an entity—static embedding  $e$  based on row  $r_i$  and



aggregated embedding  $\bar{e}$  based on all rows where the entity appears. In the context of the RotoWire dataset, the aggregate embedding  $\bar{e}$  is supposed to represent how entity  $e$  played in the game. For  $h_t = \{h_t^{lm}, h_t^{ent}\}$ ,  $P(z_t = 1|h_{t-1})$  (Equation (4)) models the transition probability, and based on whether  $e$  belongs to the set of entities  $\epsilon_t$  that have already appeared at time step  $t$ ,  $P(e_t = e|h_{t-1})$  (Equation (5)) computes the next probable entity  $e$  to mention. The authors note that such discrete tracking dramatically suppresses the generation of redundant relations in the narrative

$$P(z_t = 1|h_{t-1}) = \sigma(W_1(h_{t-1}^{lm} \oplus h_{t-1}^{ent})) \quad (4)$$

$$P(e_t = e|h_{t-1}) \propto \begin{cases} e^{(h_s^{ent} W_1 h_{t-1}^{lm})} & e \in \epsilon_{t-1} \\ e^{(\bar{e} W_2 h_{t-1}^{lm})} & \text{otherwise} \end{cases} \quad (5)$$

With the premise that paragraphs are the smallest sub-categorization where coherence and topic are defined [358], Puduppully and Lapata [255] propose a paragraph-based *macro* planning framework specific to the design of MLB [253] and RotoWire [347] datasets where the input to the seq2seq framework are predicted macro-plans (sequence of paragraphs). Building upon this, in contrast to precomputing global macro plans, Puduppully et al. [254] interweave the macro planning process with narrative generation where latent plans are sequentially inferred through a structured variational model as the narrative is generated conditioned on the plans so far and the previously generated paragraphs. Similarly, to establish order in the generation process, Sha et al. [292] incorporate link-based attention [114] in addition to content-based attention [11] into their framework. Similar to transitions in Markov chains [155], a link matrix  $\mathbb{L} \in \mathbb{R}^{n_f \times n_f}$  for  $n_f$  tabular attributes defines the likelihood of transitioning from the mention of attribute  $i$  to  $j$  as  $\mathbb{L}(f_j, f_i)$ . Wang et al. [337] propose combining autoregressive modeling [287] to generate skeletal plans and using an iterative text-editing based non-autoregressive decoder [120] to generate narratives constrained on said skeletal plans. The authors note that this approach reduces hallucination tendencies of the model. Similarly, motivated by the strong correlation observed between entity-centric metrics for record coverage and hallucinations, Liu et al. [194] adopt a two-stage generation process where a plan generator first transforms the input table records into serialized plans  $R \rightarrow R + P$  based on the separator token *SEP* and then translates the plans into narratives with the help of appended auxiliary entity information extracted through NER.

*Handcrafted templates* traditionally served as pre-defined structures where entities computed through CS would be plugged-in. However, even in the neural D2T paradigm, inducing underlying templates helps capture the narrator voicing and stylistic representations present in the training set. As such, Ye et al. [356] extend the use of the **variational auto-encoders (VAEs)** [163] for template induction with their variational template machine that disentangles the latent representation of the template  $z$  and the content  $c$ . In essence, the model can be trained to follow specific templates by sampling from  $z$ . Inspired from stylistic encoders [137], the authors further promote template learning by anonymizing entities in the input table thus effectively masking the CS process. Similarly, to mitigate the strong model biases in the standard conditional VAEs [319], Chen et al. [49] estimate semantic *confounders*  $z_c$ —linguistically similar entities to the target tokens that confound the logic of the narrative. Compared to the standard formulation  $p(y|x)$ , the authors employ Pearl’s do-calculus [242] to learn the objective  $p(y|\text{do}(x))$  that asserts that confounder  $z_c$  is no longer determined by instance  $x$ , thus ensuring logical consistency in the narrative. To ensure that the estimated confounders are meaningful, they are grounded through proxy variables  $c$  such that confounding generation  $p(c|z_m)$  can be minimized. Recently, modeling D2T and T2D as complementary tasks, Doung et al. [77] leverage the VAE formulation with the underlying architecture of a pre-trained T5 model to offer a unified multi-domain framework for the dual task. To combat the lack of

	Name	Food	Area	Price	Near
<b>Table</b>	Loch Fyne	Italian	Riverside	€20-25	Strada
<b>Exemplar <math>y_e</math></b>	<i>Zizzi is a pub providing fine French dining but with an expensive price, located near Cocum in the city center.</i>				
<b>Generation <math>y</math></b>	<i>Loch Fyne provides fine Italian dining with a £20-25 price, located near Strada at the riverside.</i>				

Fig. 9. Style imitation with exemplar narratives—Lin et al. [186].

parallel-corpora for the back-translation (T2D) training, the authors introduce latent variables to model the marginal probabilities of back-translation through an iterative learning process.

Likewise, for approaches beyond the use of auto-encoders, Chen et al. [47] take inspiration from practices in semantic parsing [71] and propose a coarse-to-fine two-stage generation scheme. In the first stage, a template  $Y_T$  containing placeholder tokens  $ENT$  is generated, representing the global logical structure of the narrative. The entities are then copied over from the input data instance to replace tokens  $ENT$  in the second step to generate the final narrative  $\hat{Y}$ . Suadaa et al. [312], similarly, follow template-guided generation [151] (see Section 4.2) where the precomputed results of numeric operations are copied over to the template and replace the placeholder tokens. For **pre-trained language models (PLMs)**, the authors incorporate copying into the fine-tuning stage for this action.

**5.1.4 Stylistic Encoders.** In addition to the traits of coherence, fluency, and fidelity, stylistic variation is crucial to NLG [307]. It is interesting to note that the n-gram entropy of generated texts in seq2seq-based NLG systems are significantly lower than that in its training data—leading to the conclusion that these systems adhere to only a handful of dominant patterns observed in the training set [237]. Thus, introducing control measures to text generation has recently garnered significant attention from the NLG community [137, 344]. As such, the semantically conditioned LSTM proposed by Wen et al. [343] extends the LSTM cell to incorporate a one-hot encoded MR vector  $d$  that takes the form of a sentence planner. Following this, Deriu and Cieliebak [62] append additional syntactic control measures to the MR vector  $d$  (such as the first token to appear in the utterances and expressions for different entity–value pairs) by simply appending one-hot vectors representation of these control mechanisms to  $d$ . Similarly, Lin et al. [186] tackle the lack of a template-based parallel dataset with *style imitation*—as illustrated in Figure 9, for each instance  $(x, y)$ , an *exemplar narrative*  $y_e$  is retrieved from the training set based on field-overlap distance  $D(x, x_e)$  and an additional encoder is used to encode  $y_e$ . The model is trained with competing objectives for content determination  $P(y|x, y_e)$  and style embodiment  $P(y_e|x_e, y_e)$  with an additional content coverage constraint for better generation fidelity.

**5.1.5 Graph Encoders.** The use of explicit graph encoders in D2T stems from the intuition that neural graph encoders such as **graph convolutional networks (GCNs)** [164] have strong relational inductive biases that produce better representations of input graphs [18] as an effective alternative to linearization. This entails generating representations for the nodes  $v \in V$  and edges  $(u, v) \in E$  in the input graph.

**GCNs and graph-RNNs:** Marcheggiani and Titov [208] compute node Representations  $h'_v$  (6) through explicit modeling of edge labels  $lab(u, v)$  and directions  $dir(u, v) \in \{in, out, loop\}$  for each neighboring node  $u \in N(v)$  in their GCN parameterization where learned scalar gates  $g_{u,v}$  weigh

the importance of each edge. With residual ( $h'_v = h'_v + h_v$ ) [129] and dense ( $h^d = [h'_v; h_v]$ ) [138] skip connections, Marcheggiani and Perez-Beltrachini [207] adopt the above-mentioned encoder with an LSTM decoder [201] for graph-to-text generation. Differing from previous iterations of Graph LSTMs [184], Distawan et al. [68] compute the hidden states of graph entities with consideration of the edges pointing to the entity from the previous entities, allowing their GTR-LSTM framework to handle non-predefined relationships. The ordering of the vertices fed into the LSTM is based on a combination of topological sort and breath-first traversal. Inspired by hybrid traversal techniques [226, 298], Ribeiro et al. [273] propose a dual graph encoder—first operating on a top-down traversal of the input graph where the predicate  $p$  between two nodes is used to transform labeled edges  $(u_i, p, u_j)$  to two unlabeled edges  $(u_i, p)$  and  $(p, u_j)$  while the second operates on a bottom-up traversal where directions of edges are reversed  $(u_i, u_j) \rightarrow (u_j, u_i)$

$$h'_v = \text{ReLU} \left( \sum_{u \in N(v)} g_{u,v} (W_{dir(u,v)} h_u + b_{lab(u,v)}) \right). \quad (6)$$

Damonte and Cohen [61] note that GCNs can assist LSTMs in capturing re-entrant structures and long term dependencies. As such, to bridge the gap between the GCN [19, 285] and the linearized LSTM encoders in graph-to-text translation, Zhao et al. [364] propose DualEnc which uses both to capture their complementary effects. The first GCN models the graph, retaining its structural integrity, while the second GCN serializes and re-orders the graph nodes resembling a planning stage and feeds it to the LSTM decoder.

*GATs and graph transformers:* To address the shortcomings of RNN-based sequential computing, Koncel-Kedziorski et al. [168] extend the transformer architecture [326] to graph-structured inputs with the GraphWriter. The distinction of GraphWriter from **graph attention networks (GATs)** [327] is made through the contextualization of each node representation  $v_i$  (7) with respect to its neighbors  $u_j \in N(v_i)$  through attention mechanism  $a_n$  for the  $N$  attention heads. In contrast, Ribeiro et al. [275] focus on capturing complementary graph contexts through distinct global  $h_v^{global}$  and local  $h_v^{local}$  message passing using GATs. Their approach to graph modeling also differs in its token-level approach for node representations with positional embeddings injected to preserve sequential order of the tokens

$$\hat{v}_i = v_i + \sum_{n=1}^N \alpha_{ij}^n W_v^n u_j \quad \alpha_{ij}^n = a^n(v_i, u_j). \quad (7)$$

Song et al. [306] enrich the training signal to the relation-aware transformer model [367] through additional multi-view auto-encoding losses [264]. This detachable multi-view framework deconstructs the input graph into triple sets for the first view, reconstructed with a deep biaffine model [73], and linearizes the graph through a depth-first traversal for the second view. In contrast, Ke et al. [158] obtain the entity and relation embeddings through contextual semantic representations with their structure-aware semantic aggregation module added to each transformer layer—the module consists of a mean pooling layer for entity and relation representations, a structure-aware self-attention layer [296], and finally a residual layer that fuses the semantic and structural representations of entities.

*5.1.6 Reconstruction and Hierarchical Decoders. Input reconstruction:* Conceptualized from auto-encoders [34, 304, 330], reconstruction-based models quantify the faithfulness of an encoded representation by correlating the decoded representation to the original input. As such, Wiseman et al. [347] adopt decoder reconstruction [321] to the D2T paradigm by segmenting the decoder hidden states  $h_t$  into  $\frac{T}{B}$  continuous blocks  $b_i$  of size at most  $B$ . The prediction of record  $r$  from such

a block  $b_i$ ,  $p(r.e.r.m|b_i)$ , is modeled as  $\text{softmax}(f(b_i))$  where  $f$  is a convolutional layer followed by an MLP. To replicate the actions of an auto-encoder, Chisholm et al. [52] train a seq2seq-based reverse re-encoding *text-to-data* model along with a forward seq2seq D2T model. Similarly, Roberti et al. [277] propose a character-level GRU implementation where the recurrent module is passed as a parameter to either the encoder or the decoder depending on the forward  $\hat{y} = f(x)$  or reverse  $\hat{x} = g(y)$  direction. Following the mechanics of back-translation (text-to-data) [290, 321], Bai et al. [12] extend the standard transformer decoder [326] to reconstruct the input graph by jointly predicting the node and edge labels while predicting the next token. The standard training objective of minimizing the negative log-likelihood of the conditional word probabilities  $l_{std}$  is appended by a node prediction loss  $l_{node}$  that minimizes the word-to-node attention distance and an edge prediction loss  $l_{edge}$  that minimizes the negative log-likelihood over the projected edges. For table-structure reconstruction, Gong et al. [109] define the reconstruction loss based on attribute prediction and content matching similar to the optimal transport distance [339]. It should be noted that these auxiliary tasks improve the model performance in few-shot settings.

*Hierarchical decoding*: Similar to hierarchical encoding (see Section 5.1.2), hierarchical *decoding* intends to designate granular roles to each decoder in the hierarchy. Serban et al. [291] show that injecting variations at the conditional output distribution does not capture high-level variations. As such, to model both high and low level variations, Shao et al. [293] propose their **planning-based hierarchical variational model (PHVM)** based on the conditional VAE [305]. PHVM follows a hierarchical multi-step encoder–decoder setup where a plan decoder first generates a subset  $g$  of the input  $\{d_i, \dots, d_n\} \in x$ . Then, in the hierarchical generation process, a *sentence decoder* and a *word decoder* generate the narrative conditioned on plan  $g$ . To dissipate the decoder responsibilities in the seq2seq paradigm, Su et al. [310] propose a four-layer hierarchical decoder where each layer is responsible for learning different parts of the output speech. The training instances are appended with POS tags such that each layer in the decoder hierarchy is responsible for decoding words associated with a specific set of POS patterns.

*Hierarchical attention-based decoding*: To alleviate omissions in narrative generation, Liu et al. [192] propose *forced attention*—with word-level coverage  $\theta_t^i$  and attribute-level coverage  $\gamma_t^e$ , a new context vector  $\hat{c}_t = \pi c_t + (1 - \pi)v_t$  is defined with a learnable vector  $\pi$  and a compensation vector  $v_t = f(\theta_t^i, \gamma_t^e)$  for low-coverage attributes  $e$ . To enforce this at a global scale, similar to Xu et al. [352], a loss function  $\mathbb{L}_{FA}$  based on  $\gamma_t^e$  is appended to the seq2seq loss function.

**5.1.7 Regularization Techniques.** Similar to regularization in the greater deep learning landscape [110], regularization practices in D2T append additional constraints to the loss function to enhance generation fidelity. As such, Mei et al. [210] introduce a coarse-to-fine aligner to the seq-to-seq framework that uses a *pre-selector* and *refiner* to modulate the standard aligner [11]. The pre-selector assigns each record a probability  $p_i$  of being selected based on which the refiner re-weights the standard aligner’s likelihood  $w_{ti}$  to  $\alpha_{ti}$ . The weighted average  $z_t = \sum_i \alpha_{ti} m_i$  is used as a soft approximation to maintain the architecture differentiability. Further, the authors regularize the model with a summation of the learned priors  $\sum_{i=1}^N p_i$  as an approximation of the number of selected records. Similarly, Perez-Beltrachini and Lapata [244] precompute binary alignment labels for each token in the output sequence indicating its alignment with some attribute in the input record. The prediction of this binary variable is used as an auxiliary training objective for the D2T model. For tabular datasets, Liu et al. [191] propose a two-level hierarchical encoder that breaks the learning of semantic tabular representation into three auxiliary tasks incorporated into the loss function of the model. The auxiliary sequence labeling task  $L_{SL}$ , learnt in unison with seq2seq learning, predicts the attribute name for each table cell. Similarly, the auto-encoder supervision  $L_{AE}$  penalizes the distance between the table  $z_t$  and the narrative  $z_b$  representations, while the

multi-label supervision task  $L_{ML}$  predicts all the attributes in the given table. The individual losses, along with the language modeling loss, defines the loss function of the framework. To mitigate information hallucination and avoid the high variance exhibited by the use of policy gradients in the reinforcement-learning paradigm, Wang et al. [339] compute two losses in addition to the language modeling loss—the first checks the disagreement between the source table and the corresponding narrative through the L2 loss between their embeddings, similar to Yang et al. [354], while the second uses optimal transport [43]-based maximum flow between the narrative and input distributions  $\mu$  and  $\nu$ . Tian et al. [318] propose the use of *confidence priors* to mitigate hallucination tendencies in T2T generation through learned confidence scores. At each decoding step  $y_t$ , instead of concatenating all the previous attention weights, only the antecedent attention weight  $a_{t-1}$  is fed back to the RNN, such that an attention score  $A_t$  can be used to compute how much  $a_t$  affects the context vector  $c_t$ —as all the source information in  $c_t$  comes from  $a_t$ . The confidence score  $C_t(y_t)$  is then used to sample target sub-sequences faithful to the source using a variational Bayes scheme [167]. Similarly, inspired by Liu et al. [191], Li et al. [180] propose two auxiliary supervision tasks incorporated into the training loss— number ranking and importance ranking, both crucial to sport summaries, modeled with pointer networks on the outputs of the row and column encoders, respectively.

**5.1.8 RL.** In the D2T premise, language-conditional RL [200] often aids in model optimization through its role as auxiliary loss functions. While traditionally, the BLEU (see Section 6.1) and TF-IDF [260] scores of generated texts were used as the basis for RL [192], Perez-Beltrachini and Lapata [244] use alignment scores of the generated text with the target text. Similarly, Gong et al. [107] use four entity-centric metrics that center around entity importance and mention. Rebuffel et al. [262] propose a model agnostic RL framework, PARENTing which uses a combination of language model loss and RL loss computed based on PARENT F-score [65] to alleviate hallucinations and omissions in T2T generation. To avoid model overfitting on weaker training samples and to ensure the rewards reflect improvement made over pretraining, self-critical training protocol [272] is applied using the REINFORCE algorithm [345]. The improvement in PARENT score over a randomly sampled candidate  $y_c$  and a baseline sequence generated using greedy decoding  $y_b$  is used as the reward policy. In contrast, Zhao et al. [365] use generative adversarial networks [111] where the generator is modeled as a policy with the current state being the generated tokens and the action defined as the next token to select. The reward for the policy is a combination of two values—the discriminator probability of the sentence being real and the correspondence between generated narrative and the input table based on the BLEU score. As RL frameworks based on singular metrics makes it difficult to simultaneously tackle the multiple facets of generation, Ghosh et al. [105] linearly combine metrics for recall, repetition, and reconstruction, along with the BLEU score, to form a composite reward function. The policy is adapted from Wang et al. [338] and trained using Maximum Entropy Inverse RL [368].

**5.1.9 Fine-tuning Pretrained Language Models.** PLMs [64, 257] have been successful in numerous text generation tasks [288, 363]. The extensive pretraining grants these models certain worldly knowledge [247] such that, at times, the models refuse to generate nonfactual narratives even when fed deliberately corrupted inputs [274]. As such, Mager et al. [203] propose an alternate approach to fine-tuning GPT-2 for AMR-to-text generation where the fine-tuning is done on the joint distribution of the AMR  $x_j$  and the text  $y_i$  as  $\prod_i^N p(y_i|y_{<i}, x_{1:M}) \cdot \prod_j^M p(x_j|x_{<j})$ . On the other hand, inspired by *task-adaptive* pretraining strategies for text classification [122], Ribeiro et al. [274] introduce supervised and unsupervised task-adaptive pretraining stages as intermediaries between the original pretraining and the fine-tuning for graph-to-text translation. Interestingly,

the authors note good performance of the task-adapted PLMs even when trained on shuffled graph representations. Chen et al. [51] note the few-shot learning capabilities of GPT-2 for T2T generation when appended with a soft switching policy for copying tokens [287]. Similarly, as a light-weight alternative to fine-tuning the entire model, Li and Liang [182] take inspiration from prompting [37], and propose *prefix-tuning* which freezes the model parameters to only optimize the prefix, a task-specific vector prepended to the input. The authors note significant improvements in low-data settings when the prefix is initialized with embeddings from task specific words such as *T2T*. For avenues that allow leveraging the worldly knowledge of PLMs even *without fine-tuning*, Xiang et al. [351] leverage a combination of prompting GPT-3 for disambiguation and T5 for sentence fusion leading to a domain-agnostic framework for D2T generation.

Inspired from practices in unlikelihood learning [323, 341], Nan et al. [223] model T5 as both a generator and a faithfulness discriminator with two additional learning objectives for unlikelihood and replacement detection. To train the model with said objectives,  $n$  contradictory sentences  $Y_{False}^{(i,j)}$  are generated for each entailed sentence  $Y_{True}^{(i)}$  wherein the discrimination probability is computed at every step of token generation. Similarly, to address omissions in D2T (Section 3.2), Jolly et al. [144] adapt the search-and-learn formulation [178] to a few-shot setting through a two-step fine-tuning process wherein a T5 model fine-tuned on the D2T task is further fine-tuned again with omitted attributes *r.e/r.m* reinserted to the narratives as pseudo-groundtruths.

*5.1.10 Supplemental Frameworks. Supplementary modules:* Fu et al. [98] propose the adaptation of the seq2seq framework for their partially aligned dataset WITA using a supportiveness adaptor and a rebalanced beam search. The pre-trained adaptor calculates supportiveness scores for each word in the generated text with respect to the input. This score is incorporated into the loss function of the seq2seq module and used to rebalance the probability distributions in the beam search. Framing the generation of narratives as a *sentence fusion* [17] task, Kasner and Dušek [157] use the pre-trained LasterTagger text editor [205] to iteratively improve a templated narrative. Su et al. [311] adopt a BM25 [278] based information retrieval system to their prototype-to-generate framework, which, aided with their BERT-based prototype selector, retrieves contextual samples for the input data instance from Wikipedia, allowing for successful few shot learning in T5. For reasoning over tabulated sport summaries, Li et al. [180] propose a variation on GATs named as GatedGAT that operates over an entity graph modeled after the source table to aid the generation model in entity-based reasoning.

*Re-ranking and pruning:* Dušek and Jurcicek [79] append the seq2seq paradigm with an RNN-based re-ranker to penalize narratives with missing and/or irrelevant attributes from the beam search output. Based on the Hamming distance between two one-hot vectors representing the presence of slot-value pairs, the classifier employs a logistic layer for a binary classification decision. Their framework, TGen, was the baseline for the E2E challenge [81]. Following this, Juraska et al. [150] first compute slot-alignment scores with a heuristic-based slot aligner which is used to augment the probability score from the seq2seq model. The aligner consists of a *gazetteer* that searches for overlapping content between the MR and its respective utterance, WordNet [87] to account for semantic relationships, and hand-crafted rules to cover the outliers. Noting that even copy-based seq2seq models tend to omit values from the input data instance, Gehrmann et al. [103] incorporate coverage  $cp$  (Equation (8)) and length  $lp$  (Equation (9)) penalties of Wu et al. [350]. With tunable parameters  $\alpha$  and  $\beta$ ,  $cp$  increases when too many generated words attend to the same input  $a_t^i$  and  $lp$  increases with the length of the generated text. In contrast to Tu et al. [322], however, the penalties are only used during inference to re-rank the beams. Similar to Paulus et al. [240], authors prune beams that start with the same bi-gram to promote syntactic variations in the generated text. Similar to **natural language inference (NLI)**-based approaches, Harkous et al.

[127] append a RoBERTa [196] based semantic fidelity classifier that reranks the beam output from a fine-tuned GPT-2 model

$$cp(x, y) = \beta \cdot \sum_{i=1}^{|x|} \log \left( \min \left( \sum_{t=1}^{|y|} a_i^t, 1 \right) \right) \quad (8)$$

$$lp(y) = \frac{(5 + |y|)^\alpha}{(5 + 1)^\alpha}. \quad (9)$$

**5.1.11 Ensemble Learning.** Juraska et al. [150] propose SLUG, an ensemble of three neural encoders—two LSTMs and one CNN [175], individually trained, for MR-text generation. The authors note that selecting tokens with the maximum log-probability by averaging over different encoders at each time steps results in incoherent narratives, thus the output candidate is selected based on the rankings among the top 10 candidates from each model. SLUG’s ensemble along with its data preprocessing and ranking schemes, as detailed in the sections above, was crowned winner of the 2017 E2E challenge [81]. Similarly, to prompt the models  $f_1, \dots, f_n$  in the ensemble to learn distinct sentence templates, Gehrmann et al. [103] adopt diverse ensembling [124] where an unobserved random variable  $w \sim \text{Cat}(1/n)$  assigns a weight to each model for each input. Constraining  $w$  to  $\{0, 1\}$  trains each model  $f_i$  on a subset of the training set thus leading to each model learning distinct templates. The final narrative is generated with a single model  $f$  with the best perplexity on the validation set.

## 5.2 Unsupervised Learning

**5.2.1 D2T Specific Pretraining.** Following the successful applications of knowledge-grounded language models [4, 197], Konstas et al. [169] propose a domain-specific pretraining strategy inspired by Sennrich et al. [290] to combat the challenges in data sparsity, wherein self-training is used to bootstrap an AMR parser from the large unlabeled Gigaword corpus [225] which is in turn used to pretrain an AMR generator. Both the generator and parser adopt the stacked-LSTM architecture [350] with a global attention decoder [201]. Similarly, following success brought forth by the suite of PLMs [37, 64, 257, 259], Chen et al. [48] propose a knowledge-grounded pretraining framework trained on 1.8 million graph-text pairs of their knowledge-grounded dataset KGTEXT built with Wikipedia hyperlinks matched to WikiData [332]. The framework consists of a graph attention network [327]-based encoder and transformer [326]-based encoders and decoders. Ke et al. [158] propose three graph-specific pretraining strategies based on the KGTEXT dataset—reconstructing masked narratives based on the input graph, conversely, reconstructing masked graph entities based on the narrative, and matching the graph and narrative embeddings with optimal transport. Similarly, Agarwal et al. [2] verbalize the entirety of the Wikidata Corpus [332] with two-step fine-tuning for T5 [259] to construct their KeLM corpus. The authors utilize this corpus to train knowledge-enhanced language models for downstream NLG tasks with significant improvements shown in the performance of REALM [123] and LAMA [247] for both retrieval as well as question answering tasks. Similarly, specifically geared for logical inference from tables, Liu et al. [187] propose PLoG wherein a PLM is first pre-trained on table-to-logic conversion intended to aid logical T2T generation for downstream datasets to the likes of LogicNLG.

**5.2.2 Auto-Encoders.** For a D2T framework solely based on unlabeled text, Freitag and Roy [95] adapt the training procedure of a denoising auto-encoder [329] to the seq2seq framework with the notion of reconstructing each training example from a partially destroyed input. For each training instance  $x_i$ , a percentage  $p$  (sampled from a Gaussian distribution) of words are removed at random to get a partially destroyed version  $\hat{x}_i$ . However, the authors note carry over of this unsupervised

approach across further D2T tasks such as the WebNLG challenge [99] could be limited by the fact that the slot names in WebNLG contribute to the MR.

### 5.3 Innovations Outside of the Seq2Seq Framework

**5.3.1 Template Induction.** For enhanced interpretability and control in D2T, Wiseman et al. [348] propose a neural parameterization of the **hidden semi-Markov model (HSMM)** [221] that jointly learns latent templates with generation. With discrete latent states  $z_t$ , length variable  $l_t$ , and a deterministic binary variable  $f_t$  that indicates whether a segment ends at time  $t$ , the HSMM is modeled as a joint-likelihood (Equation (10)). Further, associated phrases from  $x$  can be mapped to latent states  $z_t$  such that common templates can be extracted from the training dataset as a sequence of latent states  $z^i = \{z_1^i, \dots, z_S^i\}$ . Thus, the model can then be conditioned on  $z^i$  to generate text set to the template. Following this, Fu et al. [97] propose template induction by combining the expressive capacity of probabilistic models [222] with graphical models in an end-to-end fashion using a conditional random field model with Gumbel-softmax used to relax the categorical sampling process [141]. The authors note performance gains on HSMM-based baselines while also noting that neural seq2seq fare better than both

$$P(y, z, l, f|x) = \prod_{t=0}^{T-1} P(z_{t+1}, l_{t+1}|z_t, l_t, x)^{f_t} \times \prod_{t=1}^T P(y_{t-l_{t+1}:t}|z_t, l_t, x)^{f_t}. \quad (10)$$

**5.3.2 Discrete Neural Pipelines.** While the traditional D2T pipeline observes discrete modeling of the content planning and linguistic realization stages [270], neural methods consolidate these discrete steps into end-to-end learning. With their neural referring expressions generator Neural-REG [89] appended to discrete neural pipelines, and the GRU [54] and transformer [326] as base models for both, Ferreira et al. [91] compare neural implementations of these discrete pipelines to end-to-end learning. In their findings, authors note that the neural pipeline methods generalize better to unseen domains than end-to-end methods, thus alleviating hallucination tendencies. This is also corroborated by findings from Elder et al. [82].

**5.3.3 Computational Pragmatics.** Pragmatic approaches to linguistics naturally correct under-informativeness problems [117, 134] and are often employed in grounded language learning [206, 215]. Shen et al. [297] adopt the reconstructor-based [96] and distractor-based [58] models of pragmatics to MR-to-text generation. These models extend the *base speaker* models  $S_0$  using reconstructor  $R$  and distractor  $D$  based *listener* models  $L(y|x) \in \{L^R, L^D\}$  to derive *pragmatic speakers*  $S_1(y|x) \in \{S_1^R, S_1^D\}$  (11, 12) where  $\lambda$  and  $\alpha$  are rationality parameters controlling how much the model optimizes for discriminative outputs

$$S_1^R(y|x) = L^R(x|y)^\lambda \cdot S_0(y|x)^{1-\lambda} \quad (11)$$

$$S_1^D(y|x) \propto L^D(x|y)^\alpha \cdot S_0(y|x). \quad (12)$$

## 6 Evaluation of D2T Systems

In this section, we take a deeper look into the specifics of evaluation for D2T systems. Traditionally, the evaluation of D2T systems is compartmentalized into either *intrinsic* or *extrinsic* measures [26]. The former either uses automated metrics to compare the generated narrative to a reference text or employs for human judgment [146]—both evaluating the properties of the system output. The latter focuses on the ability of the D2T system to fulfill its intended purpose of imparting information—to what degree does the system achieve its overarching task for which it was developed. From the analysis of 79 papers spanning 2005–2014, Gkatzia and Mahamood [106] note the overwhelming prevalence of intrinsic evaluation with 75.7% of articles reporting it compared to 15.1% that report



an extrinsic measure. This is unsurprising, as intrinsic evaluation can be automated and is often convenient, not requiring additional crowd-sourced human labor and collection of feedback from deployed systems. As such, Reiter [267] notes the importance of extrinsic (pragmatic) evaluation and its absence in the field. Thus, the absence of literature in extrinsic evaluation measures leads us to focus on the innovations in improving the quality of intrinsic evaluation metrics (Section 6.2). For a broader view on the evaluation of text generation systems in the greater NLG landscape, we refer the readers to a recent survey of evaluation practices for text generation systems by Celikyilmaz et al. [39].

### 6.1 BLEU: The False Prophet for D2T

With the abundance of paired datasets where each data instance is accompanied by a human generated reference text, often referred to as the *gold standard*, the NLG community has sought after quick, cheap, and effective metrics for evaluation of D2T systems. The adoption of automated metrics such as BLEU, NIST, and ROUGE, by the MT community, by the virtue of their correlation with human judgment [69, 185, 238], similarly carried over to the D2T community. Among them, Belz and Gatt [22] note that NIST best correlates with human judgments on D2T texts, when compared to 9 human domain-experts and 21 non-experts. However, they note that these n-grams-based metrics perform poorer in D2T as compared to MT due to the domain-specific nature of D2T systems wherein the generated texts are judged better by humans than human-written texts.

From a review of 284 correlations reported in 34 papers, Reiter [268] notes that the correlations between BLEU and human evaluations are inconsistent—even in similar tasks. While automated metrics can aid in the diagnostic evaluation of MT systems, the author showcases the weakness of BLEU in the evaluation of D2T systems. This notion has been resonated several times [269, 286]. On top of this, undisclosed parameterization of these metrics and the variability in the tokenization and normalization schemes applied to the references can alter this score by up to 1.8 BLEU points for the same framework [250]. Similarly, it has also been shown that ROUGE tends to favor systems that produce longer summaries [313]. Further complicating the evaluation of D2T is the fact modern frameworks are neural—comparing score distributions, even with the aid of statistical significance tests, are not as meaningful due to the non-deterministic nature of neural approaches and accompanying randomized training procedures [265].

### 6.2 Innovations in Intrinsic Evaluation

Noting the shortcomings of prevalent word-overlap metrics (Section 6.1), alternative automated metrics for intrinsic evaluation have been proposed (Sections 6.2.1 and 6.2.2). To account for divergence in reference texts, Dhingra et al. [65] propose PARENT—a metric that computes precision and recall of the generated narrative  $\hat{y}$  with both the gold narrative  $y$  and its entailment to the semi-structured tabular input  $x$ .

**6.2.1 Extractive Metrics.** With dialogue generation models adopting classification-backed automated metrics [152, 179], Wiseman et al. [347] propose a relation extraction system to the likes of [60, 72] wherein the record type  $r.t$  is predicted using its corresponding entity  $r.e$  and value  $r.m$  as  $p(r.t|e, m; \theta)$ . With such a relation extraction system, the authors propose three metrics for automated evaluation:

- CS is represented by the precision and recall of unique relations extracted from  $\hat{y}_{1:t}$  that are also extracted from  $y_{1:t}$ .
- **Relation generation (RG)** is represented by the precision and number of unique relations extracted from  $\hat{y}_{1:t}$  that can be traced to  $x$ .

- **Content ordering (CO)**, similarly, by the normalized Damerau-Levenshtein distance [36] between the sequence of records extracted from  $y_{1:t}$  and  $\hat{y}_{1:t}$ .

The authors note that, given the two facets of D2T, CS pertains to *what to say* and CO to *how to say it* while RG pertains to both (factual correctness).

**6.2.2 Contextualized Metrics.** Böhm et al. [32] note that while modern frameworks for text generation compete with higher scores on automated word-overlap metrics, the quality of the generation leaves a lot to be desired. As such, the adaptation of continuous representations-based metrics shifts the focus from surface-form matching to semantic matching. Zhang et al. [362] introduce BERTScore which computes similarity scores for tokens in the system and reference text based on their BERT [64] embeddings while Mathur et al. [209] devise supervised and unsupervised metrics for NMT based on the same BERT embeddings—both having substantially higher correlation to human judgment compared to standard word-overlap metrics (see Section 6.1). Following this, Clark et al. [56] extend the word mover’s distance [173] to multi-sentence evaluation using ELMo representations [245]. Zhao et al. [366] propose the MoverScore which uses contextualized embeddings from BERT where the aggregated representations are computed based on power means [281]. Dušek and Kasner [80] employ RoBERTa [196] for NLI, where, for a given *hypothesis* and *premise*, the model computes scores for *entailment* between the two. While lower scores for forward entailment can point to omissions, backward entailment scores correspondingly point to hallucinations. Similarly, Chen et al. [47] propose parsing-based and adversarial metrics to the evaluate model correctness in logical reasoning.

**6.2.3 Human Judgment.** While human judgment is often considered to be the ultimate D2T evaluation measure, they are subject to a high degree of inconsistency (even with the same utterance), which may be attributed to the judge’s individual preferences [63, 333]—an issue that could be circumvented through a larger sample size, however such an endeavor is accompanied with an equal increase in cost for data acquisition. As such, there have been several recommendations on the proper usage of ratings and Likert scales [143, 146, 165]. From the analysis of 135 papers from specialized NLG conferences, Amidei et al. [5] note that several studies employ the Likert scale on an item-by-item basis in contrast to its design as an aggregate scale, and the analysis performed on these scales with parametric statistics do not disclose the assumptions about the distribution of the population probability. Howcroft et al. [135], on the other hand, note that the definitions of *fluency* and *accuracy* for which these scales are employed lack consistency among the papers investigated. To mitigate these inconsistencies through good experimental design, Novikova et al. [231] propose RankME, a relative-ranking based magnitude estimation method that combines the use of continuous scales [23, 113], magnitude estimation [301], and relative assessment [38]. Further, to address the quadratic growth of data required for cross-system comparisons, the authors adopt TrueSkill [132], a Bayesian data-efficient ranking algorithm used in MT evaluation [33], to RankME. The HumEval workshop [21, 24, 25] has been an invaluable resource in investigating the shortcomings and building maps to better practices in human evaluations.

### 6.3 Emphasis on Reproducibility

The last half-decade has seen the ML community place significant emphasis on the *reproducibility* of academic results [248, 302]. However, the focus of these reproducibility efforts are placed on automated metrics (Sections 6.1, 6.2.1, and 6.2.2) with the reproducibility of human evaluation results receiving far less attention. As human evaluation is often considered the ultimate measure of D2T, Belz et al. [27] initiate *ReproGen*, a shared task focused on reproducing the results of human evaluations—intended to shed better light on the reproducibility of human evaluations

and the possible interventions in design and execution of human evaluations to make them more reproducible. The authors note the inconsistency (expected) of the evaluators across different studies and hence point toward the use of metadata standardization through data-sheets such as HEDS [299]. Providing a complementary view, van der Lee et al. [325] review practices in human evaluation from 304 publications in the International Conference in NLG and the Annual Meeting of the Association of Computational Linguistics from 2018 to 2019 and outline the severe discrepancy in the spectrum of evaluator demographics and sample sizes, design practices, evaluation criteria, and put forth some common ground through a set of best practices for conducting human evaluations.

## 7 Conclusion and Future Directions

As delineated in Sections 3 and 4, innovations in D2T take inspiration from several facets of NLG and ML. From alterations to the seq2seq, pretraining-fine-tuning, auto-encoding, ensemble-learning, and reinforcement-learning paradigms, to domain-specific data preprocessing and data encoding strategies, the prospects for innovations in D2T appear as grand as that for the NLG landscape itself. Alongside, progresses made in non-anglo-centric datasets, datacards that reinforce accountability, and metrics that offer heuristic evaluation, aid in elevating D2T standards. As NLG research evolves, so will D2T, and vice versa. In the following sections, we impart our thoughts for future directions for each facet of D2T—the desiderata for D2T dataset design (Section 7.1), a forward look at the possibilities for approaches and architectures for D2T (Section 7.2) and finally, closing thoughts on the future of D2T evaluation (Section 7.3).

### 7.1 Desiderata for D2T Datasets

In Section 2, we outlined the development of parallel corpora with data-narrative pairs alongside dominant benchmark datasets in each task category. In addition to these benchmarks, it is as crucial to acknowledge niche datasets—Obeid and Hoque [233] compile a collection of 8,305 charts with their respective narratives, followed by chart-to-text [153], that encompasses 44,096 multi-domain charts. The shared gains and pitfalls in dataset design across the D2T task categories, as discussed in Section 2, offer insights that can aid the construction of future datasets with the potential to challenge the current paradigm:

- *Domain agnosticism*: Although domain-specific datasets allow the models to learn and leverage domain-specific conventions for performance gains in niche tasks, the resulting models are less malleable to unseen domains. To be adaptable and deployable for unseen niche tasks that may vary based on user requirements, it is crucial that the datasets used to train D2T models are not restricted to a single domain to avoid over-fitting on domain-specific keywords.
- *Dataset consistency*: Often, the greatest challenges for D2T systems, namely *hallucination* and *omission* (see Section 3.2), can be traced back to the datasets. Datasets facing divergence (as outlined in Section 2.3), wherein the narratives are not consistent with the data instances or vice versa, often lead to models that hallucinate or omit important aspects of the data [280].
- *Human-crafted references*: Often, to replicate human linguistics, datasets in NLP/NLG contain human annotations (narratives), considered as *gold* references. Reiter [267] notes that D2T datasets, unintentionally, may contain machine-generated annotations, such as those for WeatherGov, and urges the community to focus on human-centered narratives.
- *Linguistic diversity*: It is vital to acknowledge that the majority of the D2T benchmark datasets are anglo-centric. Joshi et al. [147] note that models built on non-anglo-centric datasets, which are fewer and far between, have the potential to impact many more people than models

built on highly resourced languages. The WebNLG 2020 challenge,<sup>10</sup> for instance, encourages submissions for both English and Russian parsing.

## 7.2 Approaches to D2T Generation: Looking Forward

In the above Sections 4 and 5, we have extensively outlined the recent innovations in D2T both inside and outside of seq2seq modeling. However, looking forward, with the emergence of highly capable **large language models (LLMs)** such as ChatGPT [235], below we discuss the reconciliation of these emergent technologies with the current D2T paradigm:

- *Few-shot learning*: D2T generation is a task that requires extrapolation beyond general linguistic understanding and commonsense reasoning, thus general LLM prompting strategies [340] may not be suited for this endeavor. Adding to that the data sparsity prevalent in D2T, extensions of prefix-tuning [182] and in-context sample search [189] may be especially favorable for building strong subset of samples for few-shot learning in LLMs.
- *Deviation from task-specific architectures*: The paradigm for D2T generation, as it stands now, prefers custom architectures, and rightfully so—they allow focused modeling of entities (Section 5.1.1) and dedicated attention mechanisms (Sections 5.1.2 and 5.1.6) to combat data infidelity that occurs as a consequence of RNN’s lack of long-form coherence. Transformer-based LLMs, however, may inadvertently model these dependencies and attention mechanisms as a function of their self-attention modules, thus allowing the convergence toward a universal architecture.
- *Effective linearization*: In line with the point above, the preference for plan-based approaches to D2T (Section 5.1.3) stems from issues in coherence. While there has been extensive work in linearizing graphs and tables (Section 4.2) showcasing that linearization in LLMs can be as effective, if not more, compared to dedicated encoders designed to capture inductive biases (Section 5.1.1, 5.1.2, and 5.1.5), recent work suggests that LLMs can handle unexpected tasks simply through their linear transcription into linguistic sentences (the LIFT framework) [67]. This line of work has extensive potential for modeling plans through their transcriptions into sentences.
- *Numeracy for D2T generation*: The NLP niche of building LLMs capable of quantitative reasoning (often referred to as Math-AI or Math-NLP) has garnered significant interest from the research community [295, 317, 334]. Although there are works in D2T that incorporate this aspect [107], the two research niches are often disparate. D2T, a field that aims to combat hallucination of data points, has a lot to gain from the advances in Math-NLP that enable models to better reason about said data points.
- *Interfacing with external APIs*: In line with the above point, interfacing LLMs with computational APIs (Wolfram Alpha [349] and ToolFormer [284]) has showcased significant enhancement potential for these already capable models. This paves a path in D2T where we deviate from viewing the generation of narratives as a sequential input-to-output mapping but rather a more involved loop comprising of numerical and logical reasoners, computational engines, pattern matchers, and validators that combine to form the greater NLG pipeline.

## 7.3 The Future of D2T Evaluation

While NLP systems generally have benchmark datasets that closely resemble the target tasks that these systems are intended to be deployed for (summarization, sentiment analysis, and language translation), D2T systems are highly specialized to the incoming data stream which differs from user to user. Thus, a one-size-fits-all approach to benchmarking, especially with automated metrics

<sup>10</sup>[https://webnlg-challenge.loria.fr/challenge\\_2020/](https://webnlg-challenge.loria.fr/challenge_2020/)

on benchmark datasets, can not showcase the utility of these systems in the real world, leading to an urgency for practical tools for extrinsic evaluation. Additionally, besides the need for fluidity and fidelity, systems placed in the real world require accountability [214]. The current leaderboard system poses the risk of blind metric optimization with disregard to model size and fairness [85]. For a holistic approach to evaluation, Gehrmann et al. [102] propose a living benchmark, GEM, similar to likes to Dynabench [161], providing challenge sets (a set of curated test sets intended to be challenging) and benchmark datasets accompanied with their D2T-specific data cards [28].

Further, as unified D2T frameworks become more decentralized with growing user-bases, the designers of these systems can utilize the user-interaction logs as measures for extrinsic evaluation, similar to the likes of ChatGPT [235]. While the D2T community places greater emphasis on measures to evaluate the *quality* of the generated narrative, the *utility* of these narratives can be evaluated with task effectiveness with and without the presence of narratives [55, 336].

## References

- [1] Rob Abbott, Brian Ecker, Pranav Anand, and Marilyn Walker. 2016. Internet Argument Corpus 2.0: An Sql Schema for Dialogic Social Media and the Corpora to go with it. In *Proceedings of the 10th International Conference on Language Resources and Evaluation (LREC'16)*. 4445–4452.
- [2] Oshin Agarwal, Heming Ge, Siamak Shakeri, and Rami Al-Rfou. 2021. Knowledge Graph Based Synthetic Corpus Generation for Knowledge-Enhanced Language Model Pre-Training. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*. 3554–3565.
- [3] Shubham Agarwal and Marc Dymetman. 2017. A Surprisingly Effective Out-of-the-Box Char2char Model on the E2E NLG Challenge Dataset. In *Proceedings of the 18th Annual SIGdial Meeting on Discourse and Dialogue*. 158–163.
- [4] Sungjin Ahn, Heeyoul Choi, Tanel Pärnamaa, and Yoshua Bengio. 2016. A Neural Knowledge Language Model. arXiv:1608.00318.
- [5] Jacopo Amidei, Paul Piwek, and Alistair Willis. 2019. The Use of Rating and Likert Scales in Natural Language Generation Human Evaluation Tasks: A Review and Some Recommendations. In *Proceedings of the 12th International Conference on Natural Language Generation*. 397–402.
- [6] Waleed Ammar, Dirk Groeneveld, Chandra Bhagavatula, Iz Beltagy, Miles Crawford, Doug Downey, Jason Dunkelberger, Ahmed Elgohary, Sergey Feldman, Vu Ha, Rodney Kinney, Sebastian Kohlmeier, Kyle Lo, Tyler Murray, Hsu-Han Ooi, Matthew Peters, Joanna Power, Sam Skjonsberg, Lucy Lu Wang, Chris Wilhelm, Zheng Yuan, Madeleine van Zuylen, and Oren Etzioni. 2018. Construction of the Literature Graph in Semantic Scholar. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, Vol. 3, Industry Papers. 84–91.
- [7] Ewa Andrejczuk, Julian Eisenschlos, Francesco Piccinno, Syrine Krichene, and Yasemin Altun. 2022. Table-To-Text Generation and Pre-Training with TabT5. In *Findings of the Association for Computational Linguistics: EMNLP 2022*. Association for Computational Linguistics, Abu Dhabi, United Arab Emirates, 6758–6766. Retrieved from <https://aclanthology.org/2022.findings-emnlp.503>
- [8] Gabor Angeli, Percy Liang, and Dan Klein. 2010. A Simple Domain-Independent Probabilistic Approach to Generation. In *Proceedings of the 2010 Conference on Empirical Methods in Natural Language Processing*. 502–512.
- [9] Luca Anselma and Alessandro Mazzei. 2018. Designing and Testing the Messages Produced by a Virtual Dietitian. In *Proceedings of the 11th International Conference on Natural Language Generation (INLG'18)*. Association for Computational Linguistics, 244–253.
- [10] Tatsuya Aoki, Akira Miyazawa, Tatsuya Ishigaki, Keiichi Goshima, Kasumi Aoki, Ichiro Kobayashi, Hiroya Takamura, and Yusuke Miyao. 2018. Generating Market Comments Referring to External Resources. In *Proceedings of the 11th International Conference on Natural Language Generation*. Association for Computational Linguistics, 135–139. DOI: <https://doi.org/10.18653/v1/W18-6515>
- [11] Dzmitry Bahdanau, KyungHyun Cho, and Yoshua Bengio. 2015. Neural Machine Translation by Jointly Learning to Align and Translate. In *Proceedings of the Third International Conference on Learning Representations (ICLR'15)*.
- [12] Xuefeng Bai, Linfeng Song, and Yue Zhang. 2020. Online Back-Parsing for AMR-to-Text Generation. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*. 1206–1219.
- [13] Laura Banarescu, Claire Bonial, Shu Cai, Madalina Georgescu, Kira Griffitt, Ulf Hermjakob, Kevin Knight, Philipp Koehn, Martha Palmer, and Nathan Schneider. 2013. Abstract Meaning Representation for Sembanking. In *Proceedings of the 7th Linguistic Annotation Workshop and Interoperability with Discourse*. 178–186.

- [14] Junwei Bao, Duyu Tang, Nan Duan, Zhao Yan, Ming Zhou, and Tiejun Zhao. 2018. Text Generation from Tables. *IEEE/ACM Transactions on Audio, Speech, and Language Processing* 27, 2 (2018), 311–320.
- [15] Regina Barzilay and Mirella Lapata. 2005. Collective Content Selection for Concept-to-Text Generation. In *Proceedings of Human Language Technology Conference and Conference on Empirical Methods in Natural Language Processing*. Association for Computational Linguistics, 331–338.
- [16] Regina Barzilay and Lillian Lee. 2004. Catching the Drift: Probabilistic Content Models, with Applications to Generation and Summarization. In *Proceedings of the Human Language Technology Conference of the North American Chapter of the Association for Computational Linguistics (HLT-NAACL'04)*. Association for Computational Linguistics, 113–120.
- [17] Regina Barzilay and Kathleen R. McKeown. 2005. Sentence Fusion for Multidocument News Summarization. *Computational Linguistics* 31, 3 (2005), 297–328.
- [18] Peter W. Battaglia, Jessica B. Hamrick, Victor Bapst, Alvaro Sanchez-Gonzalez, Vinícius Flores Zambaldi, Mateusz Malinowski, Andrea Tacchetti, David Raposo, Adam Santoro, Ryan Faulkner, Çağlar Gülçehre, H. Francis Song, Andrew J. Ballard, Justin Gilmer, George E. Dahl, Ashish Vaswani, Kelsey R. Allen, Charles Nash, Victoria Langston, Chris Dyer, Nicolas Heess, Daan Wierstra, Pushmeet Kohli, Matthew Botvinick, Oriol Vinyals, Yujia Li, and Razvan Pascanu. 2018. Relational Inductive Biases, Deep Learning, and Graph Networks. arXiv:1806.01261
- [19] Daniel Beck, Gholamreza Haffari, and Trevor Cohn. 2018. Graph-to-Sequence Learning using Gated Graph Neural Networks. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics*, Vol. 1, Long Papers. 273–283.
- [20] Anja Belz. 2008. Automatic Generation of Weather Forecast Texts using Comprehensive Probabilistic Generation-Space Models. *Natural Language Engineering* 14, 4 (2008), 431–455.
- [21] Anja Belz, Shubham Agarwal, Yvette Graham, Ehud Reiter, and Anastasia Shimorina. 2021a. Proceedings of the Workshop on Human Evaluation of NLP Systems (HumEval). In *Proceedings of the Workshop on Human Evaluation of NLP Systems (HumEval)*.
- [22] Anja Belz and Albert Gatt. 2008. Intrinsic vs. Extrinsic Evaluation Measures for Referring Expression Generation. In *Proceedings of ACL-08: HLT Short Papers*. 197–200.
- [23] Anja Belz and Eric Kow. 2011. Discrete vs. Continuous Rating Scales for Language Evaluation in NLP. In *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies*. 230–235.
- [24] Anja Belz, Maja Popović, Ehud Reiter, and Anastasia Shimorina. 2022. Proceedings of the 2nd Workshop on Human Evaluation of NLP Systems (HumEval). In *Proceedings of the 2nd Workshop on Human Evaluation of NLP Systems (HumEval)*.
- [25] Anya Belz, Maja Popović, Ehud Reiter, Craig Thomson, and João Sedoc (Eds.). 2023. *Proceedings of the 3rd Workshop on Human Evaluation of NLP Systems*. INCOMA Ltd., Shoumen, Bulgaria; Varna, Bulgaria. Retrieved from <https://aclanthology.org/2023.humeval-1.0>
- [26] Anja Belz and Ehud Reiter. 2006. Comparing Automatic and Human Evaluation of NLG Systems. In *Proceedings of the 11th Conference of the European Chapter of the Association for Computational Linguistics*. 313–320.
- [27] Anja Belz, Anastasia Shimorina, Shubham Agarwal, and Ehud Reiter. 2021b. The ReprGen Shared Task on Reproducibility of Human Evaluations in NLG: Overview and Results. In *Proceedings of the 14th International Conference on Natural Language Generation*. 249–258.
- [28] Emily M. Bender and Batya Friedman. 2018. Data Statements for Natural Language Processing: Toward Mitigating System Bias and Enabling Better Science. *Transactions of the Association for Computational Linguistics* 6 (2018), 587–604.
- [29] Yoshua Bengio, Patrice Simard, and Paolo Frasconi. 1994. Learning Long-Term Dependencies with Gradient Descent is Difficult. *IEEE Transactions on Neural Networks* 5, 2 (1994), 157–166.
- [30] Arianna Bisazza and Marcello Federico. 2016. A Survey of Word Reordering in Statistical Machine Translation: Computational Models and Language Phenomena. *Computational Linguistics* 42, 2 (2016), 163–205.
- [31] Thomas G. Moher, David C. Mak, Brad Blumenthal and Laura Marie Leventhal. 1993. Comparing the Comprehensibility of Textual and Graphical Programs: The Case of Petri Nets. In *Empirical Studies of Programmers: Fifth Workshop*. Ablex Publishing Corporation, 137–161.
- [32] Florian Böhm, Yang Gao, Christian M. Meyer, Ori Shapira, Ido Dagan, and Iryna Gurevych. 2019. Better Rewards Yield Better Summaries: Learning to Summarise Without References. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*. 3110–3120.
- [33] Ondřej Bojar, Rajen Chatterjee, Christian Federmann, Yvette Graham, Barry Haddow, Matthias Huck, Antonio Jimeno Yepes, Philipp Koehn, Varvara Logacheva, Christof Monz, Matteo Negri, Aurélie Névéal, Mariana Neves, Martin Popel, Matt Post, Raphael Rubino, Carolina Scarton, Lucia Specia, Marco Turchi, Karin Verspoor, and Marcos

- Zampieri. 2016. Findings of the 2016 Conference on Machine Translation. In *Proceedings of the First Conference on Machine Translation*, Vol. 2, Shared Task Papers. 131–198.
- [34] Hervé Bourlard and Yves Kamp. 1988. Auto-Association by Multilayer Perceptrons and Singular Value Decomposition. *Biological Cybernetics* 59, 4 (1988), 291–294.
- [35] Daniel Braun, Ehud Reiter, and Advait Siddharthan. 2018. SaferDrive: An NLG-Based Behaviour Change Support System for Drivers. *Natural Language Engineering* 24, 4 (2018), 551–588.
- [36] Eric Brill and Robert C. Moore. 2000. An Improved Error Model for Noisy Channel Spelling Correction. In *Proceedings of the 38th Annual Meeting of the Association for Computational Linguistics*. 286–293.
- [37] Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D. Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel M. Ziegler, Jeffrey Wu, Clemens Winter, Christopher Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. 2020. Language Models are Few-Shot Learners. *Advances in Neural Information Processing Systems* 33 (2020), 1877–1901.
- [38] Chris Callison-Burch, Cameron Shaw Fordyce, Philipp Koehn, Christof Monz, and Josh Schroeder. 2007. (Meta-) Evaluation of Machine Translation. In *Proceedings of the Second Workshop on Statistical Machine Translation*. 136–158.
- [39] Asli Celikyilmaz, Elizabeth Clark, and Jianfeng Gao. 2020. Evaluation of Text Generation: A Survey. arXiv:2006.14799.
- [40] Wallace Chafe. 1976. Givenness, contrastiveness, definiteness, subjects, topics, and point of view. In *Subject and Topic*. C. N. Li (Ed.), Academic Press, New York, NY, 25–56.
- [41] Sarath Chandar A. P., Stanislas Lauly, Hugo Larochelle, Mitesh Khapra, Balaraman Ravindran, Vikas C. Raykar, and Amrita Saha. 2014. An Autoencoder Approach to Learning Bilingual Word Representations. *Advances in Neural Information Processing Systems* 27 (2014).
- [42] David L. Chen and Raymond J. Mooney. 2008. Learning to Sportscast: A Test of Grounded Language Acquisition. In *Proceedings of the 25th International Conference on Machine Learning*. 128–135.
- [43] Liqun Chen, Yizhe Zhang, Ruiyi Zhang, Chenyang Tao, Zhe Gan, Haichao Zhang, Bai Li, Dinghan Shen, Changyou Chen, and Lawrence Carin. 2018. Improving Sequence-to-Sequence Learning via Optimal Transport. In *International Conference on Learning Representations*.
- [44] Miao Chen, Xinjiang Lu, Tong Xu, Yanyan Li, Jingbo Zhou, Dejing Dou, and Hui Xiong. 2022. Towards Table-to-Text Generation with Pretrained Language Model: A Table Structure Understanding and Text Deliberating Approach. In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*. Association for Computational Linguistics, Abu Dhabi, United Arab Emirates.
- [45] Mingda Chen, Sam Wiseman, and Kevin Gimpel. 2021b. WikiTableT: A Large-Scale Data-to-Text Dataset for Generating Wikipedia Article Sections. In *Findings of the Association for Computational Linguistics (ACL-IJCNLP'21)*. 193–209.
- [46] Shuang Chen, Jinpeng Wang, Xiaocheng Feng, Feng Jiang, Bing Qin, and Chin-Yew Lin. 2019b. Enhancing Neural Data-to-Text Generation Models with External Background Knowledge. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*. 3022–3032.
- [47] Wenhui Chen, Jianshu Chen, Yu Su, Zhiyu Chen, and William Yang Wang. 2020a. Logical Natural Language Generation from Open-Domain Tables. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*. 7929–7942.
- [48] Wenhui Chen, Yu Su, Xifeng Yan, and William Yang Wang. 2020c. KGPT: Knowledge-Grounded Pre-Training for Data-to-Text Generation. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*. 8635–8648.
- [49] Wenqing Chen, Jidong Tian, Yitian Li, Hao He, and Yaohui Jin. 2021a. De-Confounded Variational Encoder-Decoder for Logical Table-to-Text Generation. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing*, Vol. 1, Long Papers. 5532–5542.
- [50] Wenhui Chen, Hongmin Wang, Jianshu Chen, Yunkai Zhang, Hong Wang, Shiyang Li, Xiyu Zhou, and William Yang Wang. 2019a. Tabfact: A Large-Scale Dataset for Table-Based Fact Verification. arXiv:1909.02164.
- [51] Zhiyu Chen, Harini Eavani, Wenhui Chen, Yinyin Liu, and William Yang Wang. 2020b. Few-Shot NLG with Pre-Trained Language Model. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*. 183–190.
- [52] Andrew Chisholm, Will Radford, and Ben Hachey. 2017. Learning to Generate One-Sentence Biographies from Wikidata. In *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics*, Vol. 1, Long Papers. 633–642.
- [53] Kyunghyun Cho. 2016. Noisy Parallel Approximate Decoding for Conditional Recurrent Language Model. arXiv:1605.03835.

- [54] Kyunghyun Cho, Bart van Merriënboer, Dzmitry Bahdanau, and Yoshua Bengio. 2014. On the Properties of Neural Machine Translation: Encoder–Decoder Approaches. In *Proceedings of SSST-8, Eighth Workshop on Syntax, Semantics and Structure in Statistical Translation*. 103–111.
- [55] Arjun Choudhry, Mandar Sharma, Pramod Chundury, Thomas Kapler, Derek W. S. Gray, Naren Ramakrishnan, and Niklas Elmquist. 2020. Once Upon a Time in Visualization: Understanding the Use of Textual Narratives for Causality. *IEEE Transactions on Visualization and Computer Graphics* 27, 2 (2020), 1332–1342.
- [56] Elizabeth Clark, Asli Celikyilmaz, and Noah A. Smith. 2019. Sentence Mover’s Similarity: Automatic Evaluation for Multi-Sentence Texts. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*. 2748–2760.
- [57] Herbert H. Clark. 1977. Comprehension and the Given-New Contract. *Discourse Production and Comprehension. Discourse Processes: Advances in Research and Theory*. Roy O. Freedle (Ed.). Ablex Publishing., Norwood, NJ. 1–40.
- [58] Reuben Cohn-Gordon, Noah Goodman, and Christopher Potts. 2018. Pragmatically Informative Image Captioning with Character-Level Inference. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, Vol. 2, Short Papers. 439–443.
- [59] Emilie Colin and Claire Gardent. 2019. Generating Text from Anonymised Structures. In *Proceedings of the 12th International Conference on Natural Language Generation*. 112–117.
- [60] Ronan Collobert, Jason Weston, Léon Bottou, Michael Karlen, Koray Kavukcuoglu, and Pavel Kuksa. 2011. Natural Language Processing (Almost) from Scratch. *Journal of Machine Learning Research* 12 (2011), Article, 2493–2537.
- [61] Marco Damonte and Shay B. Cohen. 2019. Structural Neural Encoders for AMR-to-Text Generation. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, Vol. 1, Long and Short Papers. 3649–3658.
- [62] Jan Milan Deriu and Mark Cieliebak. 2018. Syntactic Manipulation for Generating More Diverse and Interesting Texts. In *11th International Conference on Natural Language Generation (INLG’18)*. Association for Computational Linguistics, 22–34.
- [63] Nina Dethlefs, Heriberto Cuayáhuítl, Helen Hastie, Verena Rieser, and Oliver Lemon. 2014. Cluster-Based Prediction of User Ratings for Stylistic Surface Realisation. In *Proceedings of the 14th Conference of the European Chapter of the Association for Computational Linguistics*. 702–711.
- [64] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. BERT: Pre-Training of Deep Bidirectional Transformers for Language Understanding. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, Vol. 1, Long and Short Papers. 4171–4186.
- [65] Bhuwan Dhingra, Manaal Faruqui, Ankur Parikh, Ming-Wei Chang, Dipanjan Das, and William Cohen. 2019. Handling Divergent Reference Texts when Evaluating Table-to-Text Generation. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*. 4884–4895.
- [66] Nicholas Diakopoulos. 2019. Automating the News. Harvard University Press.
- [67] Tuan Dinh, Yuchen Zeng, Ruisi Zhang, Ziqian Lin, Michael Gira, Shashank Rajput, Jy-yong Sohn, Dimitris Papailiopoulos, and Kangwook Lee. 2022. LIFT: Language-Interfaced Fine-Tuning for Non-language Machine Learning Tasks. In *Advances in Neural Information Processing Systems*.
- [68] Bayu Distiawan, Jianzhong Qi, Rui Zhang, and Wei Wang. 2018. GTR-LSTM: A Triple Encoder for Sentence Generation from RDF Data. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics*, Vol. 1, Long Papers. 1627–1637.
- [69] George Doddington. 2002. Automatic Evaluation of Machine Translation Quality using n-Gram Co-Occurrence Statistics. In *Proceedings of the Second International Conference on Human Language Technology Research*. 138–145.
- [70] Chenhe Dong, Yinghui Li, Haifan Gong, Miaoxin Chen, Junxin Li, Ying Shen, and Min Yang. 2021. A Survey of Natural Language Generation. arXiv:2112.11739.
- [71] Li Dong and Mirella Lapata. 2018. Coarse-to-Fine Decoding for Neural Semantic Parsing. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics*, Vol. 1, Long Papers. 731–742.
- [72] Cicero dos Santos, Bing Xiang, and Bowen Zhou. 2015. Classifying Relations by Ranking with Convolutional Neural Networks. In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing*, Vol. 1, Long Papers. 626–634.
- [73] Timothy Dozat and Christopher D Manning. 2017. Deep Biaffine Attention for Neural Dependency Parsing. In *Proceedings of the 5th International Conference on Learning Representations; 5th International Conference on Learning Representations*.
- [74] Nan Duan, Duyu Tang, Peng Chen, and Ming Zhou. 2017. Question Generation for Question Answering. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*. Association for Computational Linguistics, 866–874. DOI: <https://doi.org/10.18653/v1/D17-1090>
- [75] Pablo Ariel Duboue and Kathleen R. McKeown. 2003. Statistical Acquisition of Content Selection Rules for Natural Language Generation. In *Proceedings of the 2003 Conference on Empirical Methods in Natural Language Processing*. 121–128.



- [76] Sebastian Duerr and Peter A. Gloor. 2021. Persuasive Natural Language Generation—A Literature Review. arXiv:2101.05786.
- [77] Song Duong, Alberto Lumbreras, Mike Gartrell, and Patrick Gallinari. 2023. Learning from Multiple Sources for Data-to-Text and Text-to-Data. In *International Conference on Artificial Intelligence and Statistics (AISTATS)*.
- [78] Ondřej Dušek, David M. Howcroft, and Verena Rieser. 2019. Semantic Noise Matters for Neural Natural Language Generation. In *Proceedings of the 12th International Conference on Natural Language Generation*. 421–426.
- [79] Ondřej Dušek and Filip Jurcicek. 2016. Sequence-to-Sequence Generation for Spoken Dialogue via Deep Syntax Trees and Strings. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics*, Vol. 2, Short Papers. 45–51.
- [80] Ondřej Dušek and Zdeněk Kasner. 2020. Evaluating Semantic Accuracy of Data-to-Text Generation with Natural Language Inference. In *Proceedings of the 13th International Conference on Natural Language Generation*. 131–137.
- [81] Ondřej Dušek, Jekaterina Novikova, and Verena Rieser. 2018. Findings of the E2E NLG Challenge. In *Proceedings of the 11th International Conference on Natural Language Generation*. 322–328.
- [82] Henry Elder, Jennifer Foster, James Barry, and Alexander O'Connor. 2019. Designing a Symbolic Intermediate Representation for Neural Surface Realization. In *Proceedings of the Workshop on Methods for Optimizing and Evaluating Neural Language Generation*. 65–73.
- [83] Henry Elder and Chris Hokamp. 2018. Generating High-Quality Surface Realizations Using Data Augmentation and Factored Sequence Models. In *Proceedings of the First Workshop on Multilingual Surface Realisation*. Association for Computational Linguistics, 49–53. DOI: <https://doi.org/10.18653/v1/W18-3606>
- [84] Mihail Eric, Rahul Goel, Shachi Paul, Abhishek Sethi, Sanchit Agarwal, Shuyang Gao, Adarsh Kumar, Anuj Goyal, Peter Ku, and Dilek Hakkani-Tur. 2020. MultiWOZ 2.1: A Consolidated Multi-Domain Dialogue Dataset with State Corrections and State Tracking Baselines. In *Proceedings of the 12th Language Resources and Evaluation Conference*. 422–428.
- [85] Kawin Ethayarajh and Dan Jurafsky. 2020. Utility is in the Eye of the User: A Critique of NLP Leaderboards. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*. 4846–4853.
- [86] Angela Fan, Claire Gardent, Chloé Braud, and Antoine Bordes. 2019. Using Local Knowledge Graph Construction to Scale Seq2Seq Models to Multi-Document Inputs. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and 9th International Joint Conference on Natural Language Processing*.
- [87] Christiane Fellbaum. 1998. A Semantic Network of English: The Mother of All WordNets. In *EuroWordNet: A Multilingual Database with Lexical Semantic Networks*. Springer, 137–148.
- [88] Thiago Castro Ferreira, Iacer Calixto, Sander Wubben, and Emiel Kraahmer. 2017. Linguistic Realisation as Machine Translation: Comparing Different MT Models for AMR-to-Text Generation. In *Proceedings of the 10th International Conference on Natural Language Generation*. 1–10.
- [89] Thiago Castro Ferreira, Diego Moussallem, Ákos Kádár, Sander Wubben, and Emiel Kraahmer. 2018a. NeuralREG: An End-to-End Approach to Referring Expression Generation. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics*, Vol. 1, Long Papers. 1959–1969.
- [90] Thiago Castro Ferreira, Diego Moussallem, Emiel Kraahmer, and Sander Wubben. 2018b. Enriching the WebNLG Corpus. In *Proceedings of the 11th International Conference on Natural Language Generation*. 171–176.
- [91] Thiago Castro Ferreira, Chris van der Lee, Emiel Van Miltenburg, and Emiel Kraahmer. 2019. Neural Data-to-Text Generation: A Comparison Between Pipeline and End-to-End Architectures. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*. 552–562.
- [92] Thiago Castro Ferreira, Helena Vaz, Brian Davis, and Adriana Pagano. 2021. Enriching the E2E Dataset. In *Proceedings of the 14th International Conference on Natural Language Generation*. 177–183.
- [93] Jenny Rose Finkel, Trond Grenager, and Christopher D. Manning. 2005. Incorporating Non-Local Information into Information Extraction Systems by Gibbs Sampling. In *Proceedings of the 43rd annual meeting of the association for computational linguistics (ACL'05)*. 363–370.
- [94] Mary Ellen Foster. 2019. Natural Language Generation for Social Robotics: Opportunities and Challenges. *Philosophical Transactions of the Royal Society B* 374, 1771 (2019), 20180027.
- [95] Markus Freitag and Scott Roy. 2018. Unsupervised Natural Language Generation with Denoising Autoencoders. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*. 3922–3929.
- [96] Daniel Fried, Jacob Andreas, and Dan Klein. 2018. Unified Pragmatic Models for Generating and Following Instructions. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, Vol. 1, Long Papers. 1951–1963.
- [97] Yao Fu, Chuanqi Tan, Bin Bi, Moshua Chen, Yansong Feng, and Alexander Rush. 2020b. Latent Template Induction with Gumbel-CRFs. *Advances in Neural Information Processing Systems* 33 (2020), 20259–20271.

- [98] Zihao Fu, Bei Shi, Wai Lam, Lidong Bing, and Zhiyuan Liu. 2020a. Partially-Aligned Data-to-Text Generation with Distant Supervision. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*. 9183–9193.
- [99] Claire Gardent, Anastasia Shimorina, Shashi Narayan, and Laura Perez-Beltrachini. 2017. Creating Training Corpora for NLG Micro-Planners. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics*, Vol. 1, Long Papers. Association for Computational Linguistics, 179–188. DOI: <https://doi.org/10.18653/v1/P17-1017>
- [100] Albert Gatt and Emiel Kraahmer. 2018. Survey of the State of the Art in Natural Language Generation: Core Tasks, Applications and Evaluation. *Journal of Artificial Intelligence Research* 61 (2018), 65–170.
- [101] Ruifang Ge and Raymond Mooney. 2005. A Statistical Semantic Parser that Integrates Syntax and Semantics. In *Proceedings of the Ninth Conference on Computational Natural Language Learning (CoNLL'05)*. 9–16.
- [102] Sebastian Gehrmann, Tosin Adewumi, Karmanya Aggarwal, Pawan Sasanka Ammanamanchi, Anuluwapo Aremu, Antoine Bosselut, Khyathi Raghavi Chandu, Miruna-Adriana Clinciu, Dipanjan Das, Kaustubh Dhole, Wanyu Du, Esin Durmus, Ondřej Dušek, Chris Chinenye Emezue, Varun Gangal, Cristina Garbacea, Tatsunori Hashimoto, Yufang Hou, Yacine Jernite, Harsh Jhamtani, Yangfeng Ji, Shailza Jolly, Mihir Kale, Dhruv Kumar, Faisal Ladhak, Aman Madaan, Mounica Maddela, Khyati Mahajan, Saad Mahamood, Bodhisattwa Prasad Majumder, Pedro Henrique Martins, Angelina McMillan-Major, Simon Mille, Emiel van Miltenburg, Moin Nadeem, Shashi Narayan, Vitaly Nikolaev, Andre Niyongabo Rubungo, Salomey Osei, Ankur Parikh, Laura Perez-Beltrachini, Niranjana Ramesh Rao, Vikas Raunak, Juan Diego Rodriguez, Sashank Santhanam, João Sedoc, Thibault Sellam, Samira Shaikh, Anastasia Shimorina, Marco Antonio Sobrevilla Cabezudo, Hendrik Strobelt, Nishant Subramani, Wei Xu, Diyi Yang, Akhila Yerukola, and Jiawei Zhou. 2021. The GEM Benchmark: Natural Language Generation, its Evaluation and Metrics. In *Proceedings of the 1st Workshop on Natural Language Generation, Evaluation, and Metrics (GEM'21)*. 96–120.
- [103] Sebastian Gehrmann, Falcon Dai, Henry Elder, and Alexander M. Rush. 2018. End-to-End Content and Plan Selection for Data-to-Text Generation. In *Proceedings of the 11th International Conference on Natural Language Generation*. 46–56.
- [104] Nahum Gershon and Ward Page. 2001. What Storytelling can do for Information Visualization. *Communications of the ACM* 44, 8 (2001), 31–37.
- [105] Sayan Ghosh, Zheng Qi, Snigdha Chaturvedi, and Shashank Srivastava. 2021. How Helpful is Inverse Reinforcement Learning for Table-to-Text Generation?. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing*, Vol. 2, Short Papers. 71–79.
- [106] Dimitra Gkatzia and Saad Mahamood. 2015. A Snapshot of NLG Evaluation Practices 2005-2014. In *Proceedings of the 15th European Workshop on Natural Language Generation (ENLG)*. 57–60.
- [107] Heng Gong, Wei Bi, Xiaocheng Feng, Bing Qin, Xiaojiang Liu, and Ting Liu. 2020a. Enhancing Content Planning for Table-to-Text Generation with Data Understanding and Verification. In *Findings of the Association for Computational Linguistics: EMNLP'20*. 2905–2914.
- [108] Heng Gong, Xiaocheng Feng, Bing Qin, and Ting Liu. 2019. Table-to-Text Generation with Effective Hierarchical Encoder on Three Dimensions (Row, Column and Time). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*. 3143–3152.
- [109] Heng Gong, Yawei Sun, Xiaocheng Feng, Bing Qin, Wei Bi, Xiaojiang Liu, and Ting Liu. 2020b. Tablegpt: Few-Shot Table-to-Text Generation with Table Structure Reconstruction and Content Matching. In *Proceedings of the 28th International Conference on Computational Linguistics*. 1978–1988.
- [110] Ian Goodfellow, Y. Bengio, and A. Courville. 2016. Regularization for Deep Learning. *Deep Learning*. MIT Press, 216–261.
- [111] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. 2014. Generative Adversarial Nets. *Advances in Neural Information Processing Systems* 27 (2014).
- [112] Raghav Goyal, Marc Dymetman, and Eric Gaussier. 2016. Natural Language Generation through Character-based RNNs with Finite-state Prior Knowledge. In *Proceedings of the 26th International Conference on Computational Linguistics (COLING'16)*. 1083–1092.
- [113] Yvette Graham, Timothy Baldwin, Alistair Moffat, and Justin Zobel. 2013. Continuous Measurement Scales in Human Evaluation of Machine Translation. In *Proceedings of the 7th Linguistic Annotation Workshop and Interoperability with Discourse*. 33–41.
- [114] Alex Graves, Greg Wayne, Malcolm Reynolds, Tim Harley, Ivo Danihelka, Agnieszka Grabska-Barwińska, Sergio Gómez Colmenarejo, Edward Grefenstette, Tiago Ramalho, John Agapiou, Adrià Puigdomènech Badia, Karl Moritz Hermann, Yori Zwols, Georg Ostrovski, Adam Cain, Helen King, Christopher Summerfield, Phil Blunsom, Koray Kavukcuoglu, and Demis Hassabis. 2016. Hybrid Computing Using a Neural Network with Dynamic External Memory. *Nature* 538, 7626 (2016), 471–476.
- [115] Thomas R. G. Green and Marian Petre. 1992. When Visual Programs are Harder to Read than Textual Programs. In *Human-Computer Interaction: Tasks and Organisation, Proceedings ECCE-6 (6th European Conference Cognitive Ergonomics)*, Vol. 57. Citeseer.

- [116] Thomas R. G. Green, Marian Petre, and Rachel K. E. Bellamy. 1991. *Comprehensibility of Visual and Textual Programs: A Test of Superlativism Against the 'Match-Mismatch' Conjecture*. Open University, Computer Assisted Learning Research Group.
- [117] Herbert P. Grice. 1975. Logic and conversation. In *Speech Acts*. Brill, 41–58.
- [118] Barbara J. Grosz, Scott Weinstein, and Aravind K. Joshi. 1995. Centering: A Framework for Modeling the Local Coherence of Discourse. *Computational Linguistics* 21, 2 (1995), 203–225.
- [119] Jiatao Gu, Zhengdong Lu, Hang Li, and Victor OK Li. 2016. Incorporating Copying Mechanism in Sequence-to-Sequence Learning. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics*, Vol. 1, Long Papers. 1631–1640.
- [120] Jiatao Gu, Changhan Wang, and Junbo Zhao. 2019. Levenshtein Transformer. *Advances in Neural Information Processing Systems* 32 (2019).
- [121] Çağlar Gulçehre, Sungjin Ahn, Ramesh Nallapati, Bowen Zhou, and Yoshua Bengio. 2016. Pointing the Unknown Words. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics*, Vol. 1, Long Papers. 140–149.
- [122] Suchin Gururangan, Ana Marasović, Swabha Swayamdipta, Kyle Lo, Iz Beltagy, Doug Downey, and Noah A. Smith. 2020. Don't Stop Pretraining: Adapt Language Models to Domains and Tasks. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*. 8342–8360.
- [123] Kelvin Guu, Kenton Lee, Zora Tung, Panupong Pasupat, and Mingwei Chang. 2020. Retrieval Augmented Language Model Pre-Training. In *Proceedings of the 37th International Conference on Machine Learning* (Proceedings of Machine Learning Research, Vol. 119. Hal Daumé III and Aarti Singh (Eds.)). PMLR, 3929–3938.
- [124] Abner Guzman-Rivera, Dhruv Batra, and Pushmeet Kohli. 2012. Multiple Choice Learning: Learning to Produce Multiple Structured Outputs. *Advances in Neural Information Processing Systems* 25 (2012).
- [125] Michael Alexander Kirkwood Halliday and Ruqaiya Hasan. 2014. *Cohesion in English*. Routledge.
- [126] Shuang Hao, Wenfeng Han, Tao Jiang, Yiping Li, Haonan Wu, Chunlin Zhong, Zhangjun Zhou, and He Tang. 2024. Synthetic Data in AI: Challenges, Applications, and Ethical Implications. arXiv:2401.01629.
- [127] Hamza Harkous, Isabel Groves, and Amir Saffari. 2020. Have Your Text and Use It Too! End-to-End Neural Data-to-Text Generation with Semantic Fidelity. In *Proceedings of the 28th International Conference on Computational Linguistics*. 2410–2424.
- [128] Tatsunori B. Hashimoto, Hugh Zhang, and Percy Liang. 2019. Unifying Human and Statistical Evaluation for Natural Language Generation. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, Vol. 1, Long and Short Papers. 1689–1701.
- [129] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. 2016. Deep Residual Learning for Image Recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 770–778.
- [130] Shizhu He, Cao Liu, Kang Liu, and Jun Zhao. 2017. Generating Natural Answers by Incorporating Copying and Retrieving Mechanisms in Sequence-to-Sequence Learning. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics*, Vol. 1, Long Papers. 199–208.
- [131] Xiaodong He and Li Deng. 2017. Deep Learning for Image-to-Text Generation: A Technical Overview. *IEEE Signal Processing Magazine* 34, 6 (2017), 109–116. DOI: <https://doi.org/10.1109/MSP.2017.2741510>
- [132] Ralf Herbrich, Tom Minka, and Thore Graepel. 2006. TrueSkill™: A Bayesian Skill Rating System. *Advances in Neural Information Processing Systems* 19 (2006).
- [133] Sepp Hochreiter and Jürgen Schmidhuber. 1997. Long Short-Term Memory. *Neural Computation* 9, 8 (1997), 1735–1780.
- [134] Laurence Horn. 1984. Toward a New Taxonomy for Pragmatic Inference: Q-Based and R-Based Implicature. *Meaning, Form, and Use in Context: Linguistic Applications* 11 (1984), 42.
- [135] David M. Howcroft, Anja Belz, Miruna-Adriana Clinciu, Dimitra Gkatzia, Sadid A. Hasan, Saad Mahamood, Simon Mille, Emiel Van Miltenburg, Sashank Santhanam, and Verena Rieser. 2020. Twenty Years of Confusion in Human Evaluation: NLG needs Evaluation Sheets and Standardised Definitions. In *Proceedings of the 13th International Conference on Natural Language Generation*. 169–182.
- [136] David M. Howcroft, Dietrich Klakow, and Vera Demberg. 2017. The Extended SPaRky Restaurant Corpus: Designing a Corpus with Variable Information Density. In *INTERSPEECH*. 3757–3761.
- [137] Zhiting Hu, Zichao Yang, Xiaodan Liang, Ruslan Salakhutdinov, and Eric P Xing. 2017. Toward Controlled Generation of Text. In *International Conference on Machine Learning*. PMLR, 1587–1596.
- [138] Gao Huang, Zhuang Liu, Laurens Van Der Maaten, and Kilian Q. Weinberger. 2017. Densely Connected Convolutional Networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 4700–4708.
- [139] Hayate Iso, Yui Uehara, Tatsuya Ishigaki, Hiroshi Noji, Eiji Aramaki, Ichiro Kobayashi, Yusuke Miyao, Naoaki Okazaki, and Hiroya Takamura. 2020. Learning to Select, Track, and Generate for Data-to-Text. *Journal of Natural Language Processing* 27, 3 (2020), 599–626.

- [140] Glorianna Jagfeld, Sabrina Jenne, and Ngoc Thang Vu. 2018. Sequence-to-Sequence Models for Data-to-Text Natural Language Generation: Word-vs. Character-based Processing and Output Diversity. In *Proceedings of the 11th International Conference on Natural Language Generation*. 221–232.
- [141] Eric Jang, Shixiang Gu, and Ben Poole. 2017. Categorical Reparametrization with Gumble-Softmax. In *International Conference on Learning Representations (ICLR'17)*.
- [142] Ziwei Ji, Nayeon Lee, Rita Frieske, Tiezheng Yu, Dan Su, Yan Xu, Etsuko Ishii, Yejin Bang, Andrea Madotto, and Pascale Fung. 2022. Survey of Hallucination in Natural Language Generation. arXiv:2202.03629 (2022).
- [143] Robert L. Johnson and Grant B. Morgan. 2016. *Survey Scales: A Guide to Development, Analysis, and Reporting*. Guilford Publications.
- [144] Shailza Jolly, Zi Xuan Zhang, Andreas Dengel, and Lili Mou. 2022. Search and Learn: Improving Semantic Coverage for Data-to-Text Generation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 36. 10858–10866.
- [145] Bevan Jones, Jacob Andreas, Daniel Bauer, Karl Moritz Hermann, and Kevin Knight. 2012. Semantics-Based Machine Translation with Hyperedge Replacement Grammars. In *Proceedings of COLING 2012*. 1359–1376.
- [146] Ankur Joshi, Saket Kale, Satish Chandel, and D. Kumar Pal. 2015. Likert Scale: Explored and Explained. *British Journal of Applied Science & Technology* 7, 4 (2015), 396.
- [147] Pratik Joshi, Sebastin Santy, Amar Budhiraja, Kalika Bali, and Monojit Choudhury. 2020. The State and Fate of Linguistic Diversity and Inclusion in the NLP World. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*. 6282–6293.
- [148] Dan Jurafsky and James H Martin. 2014. *Speech and Language Processing*, Vol. 3. Prentice Hall.
- [149] Juraj Juraska, Kevin Bowden, and Marilyn Walker. 2019. ViGGO: A Video Game Corpus for Data-To-Text Generation in Open-Domain Conversation. In *Proceedings of the 12th International Conference on Natural Language Generation*. 164–172.
- [150] Juraj Juraska, Panagiotis Karagiannis, Kevin Bowden, and Marilyn Walker. 2018. A Deep Ensemble Model with Slot Alignment for Sequence-to-Sequence Natural Language Generation. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, Vol. 1, Long Papers. 152–162.
- [151] Mihir Kale and Abhinav Rastogi. 2020. Template Guided Text Generation for Task-Oriented Dialogue. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*. 6505–6520.
- [152] Anjali Kannan and Oriol Vinyals. 2016. Adversarial Evaluation of Dialogue Models. In: *Workshop on Adversarial Training at Neural Information Processing Systems*.
- [153] Shankar Kanthara, Rixie Tiffany Ko Leong, Xiang Lin, Ahmed Masry, Megh Thakkar, Enamul Hoque, and Shafiq Joty. 2022. Chart-to-Text: A Large-Scale Benchmark for Chart Summarization. arXiv:2203.06486.
- [154] Nikiforos Karamanis, Massimo Poesio, Chris Mellish, and Jon Oberlander. 2004. Evaluating Centering-Based Metrics of Coherence. In *Proceedings of the 42nd Annual Meeting of the Association for Computational Linguistics (ACL'04)*. 391–398.
- [155] Samuel Karlin. 2014. *A First Course in Stochastic Processes*. Academic press.
- [156] Lauri Karttunen. 1976. Discourse Referents. In *Notes from the Linguistic Underground*. Brill, 363–385.
- [157] Zdeněk Kasner and Ondřej Dušek. 2020. Data-to-Text Generation with Iterative Text Editing. In *Proceedings of the 13th International Conference on Natural Language Generation*. 60–67.
- [158] Pei Ke, Haozhe Ji, Yu Ran, Xin Cui, Liwei Wang, Linfeng Song, Xiaoyan Zhu, and Minlie Huang. 2021. JointGT: Graph-Text Joint Representation Learning for Text Generation from Knowledge Graphs. In *Findings of the Association for Computational Linguistics: ACL-IJCNLP'21*. 2526–2538.
- [159] Chris Kedzie and Kathleen Mckeown. 2019. A Good Sample is Hard to Find: Noise Injection Sampling and Self-Training for Neural Language Generation Models. In *Proceedings of the 12th International Conference on Natural Language Generation*. 584–593.
- [160] Chloé Kiddon, Luke Zettlemoyer, and Yejin Choi. 2016. Globally Coherent Text Generation with Neural Checklist Models. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*. 329–339.
- [161] Douwe Kiela, Max Bartolo, Yixin Nie, Divyansh Kaushik, Atticus Geiger, Zhengxuan Wu, Bertie Vidgen, Grusha Prasad, Amanpreet Singh, Pratik Ringshia, Zhiyi Ma, Tristan Thrush, Sebastian Riedel, Zeerak Waseem, Pontus Stenertorp, Robin Jia, Mohit Bansal, Christopher Potts, and Adina Williams. 2021. Dynabench: Rethinking Benchmarking in NLP. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*. 4110–4124.
- [162] Yoon Kim. 2014. Convolutional Neural Networks for Sentence Classification. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*. Association for Computational Linguistics, 1746–1751. DOI: <https://doi.org/10.3115/v1/D14-1181>
- [163] Diederik P. Kingma and Max Welling. 2014. Stochastic Gradient VB and the Variational Auto-Encoder. In *Second International Conference on Learning Representations (ICLR)*, Vol. 19. 121.

- [164] Thomas N. Kipf and Max Welling. 2016. Semi-Supervised Classification with Graph Convolutional Networks. In *International Conference on Learning Representations*.
- [165] Thomas R. Knapp. 1990. Treating Ordinal Scales as Interval Scales: An Attempt to Resolve the Controversy. *Nursing Research* 39, 2 (1990), 121–123.
- [166] Philipp Koehn, Hieu Hoang, Alexandra Birch, Chris Callison-Burch, Marcello Federico, Nicola Bertoldi, Brooke Cowan, Wade Shen, Christine Moran, Richard Zens, Chris Dyer, Ondřej Bojar, Alexandra Constantin, and Evan Herbst. 2007. Moses: Open source toolkit for statistical machine translation. In *Proceedings of the 45th Annual Meeting of the Association for Computational Linguistics Companion Volume Proceedings of the Demo and Poster Sessions*. 177–180.
- [167] Daphne Koller and Nir Friedman. 2009. *Probabilistic Graphical Models: Principles and Techniques*. MIT press.
- [168] Rik Koncel-Kedziorski, Dhanush Bekal, Yi Luan, Mirella Lapata, and Hannaneh Hajishirzi. 2019. Text Generation from Knowledge Graphs with Graph Transformers. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, Vol. 1, Long and Short Papers. 2284–2293.
- [169] Ioannis Konstas, Srinivasan Iyer, Mark Yatskar, Yejin Choi, and Luke Zettlemoyer. 2017. Neural AMR: Sequence-to-Sequence Models for Parsing and Generation. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics*, Vol. 1, Long Papers. 146–157.
- [170] Ioannis Konstas and Mirella Lapata. 2012. Unsupervised Concept-to-Text Generation with Hypergraphs. In *Proceedings of the 2012 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*. 752–761.
- [171] Ioannis Konstas and Mirella Lapata. 2013. Inducing Document Plans for Concept-to-Text Generation. In *Proceedings of the 2013 Conference on Empirical Methods in Natural Language Processing*. Association for Computational Linguistics, 1503–1514.
- [172] Susumu Kuno. 1972. Functional Sentence Perspective: A Case Study from Japanese and English. *Linguistic Inquiry* 3, 3 (1972), 269–320.
- [173] Matt Kusner, Yu Sun, Nicholas Kolkin, and Kilian Weinberger. 2015. From Word Embeddings to Document Distances. In *International Conference on Machine Learning*. PMLR, 957–966.
- [174] Rémi Lebret, David Grangier, and Michael Auli. 2016. Neural Text Generation from Structured Data with Application to the Biography Domain. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*. 1203–1213.
- [175] Yann LeCun, Léon Bottou, Yoshua Bengio, and Patrick Haffner. 1998. Gradient-Based Learning Applied to Document Recognition. *Proceedings of the IEEE* 86, 11 (1998), 2278–2324.
- [176] Leo Leppänen, Myriam Munezero, Mark Granroth-Wilding, and Hannu Toivonen. 2017. Data-Driven News Generation for Automated Journalism. In *Proceedings of the 10th International Conference on Natural Language Generation*. 188–197.
- [177] Uri Lerner and Slav Petrov. 2013. Source-Side Classifier Preordering for Machine Translation. In *Proceedings of the 2013 Conference on Empirical Methods in Natural Language Processing*. 513–523.
- [178] Jingjing Li, Zichao Li, Lili Mou, Xin Jiang, Michael Lyu, and Irwin King. 2020. Unsupervised Text Generation by Learning from Search. *Advances in Neural Information Processing Systems* 33 (2020), 10820–10831.
- [179] Jiwei Li, Will Monroe, Tianlin Shi, Sébastien Jean, Alan Ritter, and Dan Jurafsky. 2017. Adversarial Learning for Neural Dialogue Generation. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*.
- [180] Liang Li, Can Ma, Yinliang Yue, and Dayong Hu. 2021. Improving Encoder by Auxiliary Supervision Tasks for Table-to-Text Generation. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing*, Vol. 1, Long Papers. 5979–5989.
- [181] Wei Li, Wenhao Wu, Moye Chen, Jiachen Liu, Xinyan Xiao, and Hua Wu. 2022. Faithfulness in Natural Language Generation: A Systematic Survey of Analysis, Evaluation and Optimization Methods. arXiv:2203.05227.
- [182] Xiang Lisa Li and Percy Liang. 2021. Prefix-Tuning: Optimizing Continuous Prompts for Generation. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing*, Vol. 1, Long Papers. 4582–4597.
- [183] Percy Liang, Michael I Jordan, and Dan Klein. 2009. Learning Semantic Correspondences with Less Supervision. In *Proceedings of the Joint Conference of the 47th Annual Meeting of the ACL and the 4th International Joint Conference on Natural Language Processing of the AFNLP*. 91–99.
- [184] Xiaodan Liang, Xiaohui Shen, Jiashi Feng, Liang Lin, and Shuicheng Yan. 2016. Semantic Object Parsing with Graph lstm. In *European Conference on Computer Vision*. Springer, 125–143.
- [185] Chin-Yew Lin and Eduard Hovy. 2003. Automatic Evaluation of Summaries using n-Gram Co-Occurrence Statistics. In *Proceedings of the 2003 Human Language Technology Conference of the North American Chapter of the Association for Computational Linguistics*. 150–157.

- [186] Shuai Lin, Wentao Wang, Zichao Yang, Xiaodan Liang, Frank F Xu, Eric Xing, and Zhiting Hu. 2020. Data-to-Text Generation with Style Imitation. In *Findings of the Association for Computational Linguistics: EMNLP'20*. 1589–1598.
- [187] Ao Liu, Haoyu Dong, Naoaki Okazaki, Shi Han, and Dongmei Zhang. 2022. PLOG: Table-to-Logic Pretraining for Logical Table-to-Text Generation. In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*. Association for Computational Linguistics, Abu Dhabi, United Arab Emirates, 5531–5546. Retrieved from <https://aclanthology.org/2022.emnlp-main.373>
- [188] Fei Liu, Jeffrey Flanigan, Sam Thomson, Norman Sadeh, and Noah A. Smith. 2015. Toward Abstractive Summarization Using Semantic Representations. In *Proceedings of the 2015 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*. 1077–1086.
- [189] Jiachang Liu, Dinghan Shen, Yizhe Zhang, Bill Dolan, Lawrence Carin, and Weizhu Chen. 2021a. What Makes Good In-Context Examples for GPT-3? arXiv:2101.06804.
- [190] Shuman Liu, Hongshen Chen, Zhaochun Ren, Yang Feng, Qun Liu, and Dawei Yin. 2018a. Knowledge Diffusion for Neural Dialogue Generation. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics*, Vol. 1, Long Papers. 1489–1498.
- [191] Tianyu Liu, Fuli Luo, Qiaolin Xia, Shuming Ma, Baobao Chang, and Zhifang Sui. 2019a. Hierarchical Encoder with Auxiliary Supervision for Neural Table-to-Text Generation: Learning Better Representation for Tables. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 33. 6786–6793.
- [192] Tianyu Liu, Fuli Luo, Pengcheng Yang, Wei Wu, Baobao Chang, and Zhifang Sui. 2019b. Towards comprehensive description generation from factual attribute-value tables. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*. 5985–5996.
- [193] Tianyu Liu, Kexiang Wang, Lei Sha, Baobao Chang, and Zhifang Sui. 2018b. Table-to-Text Generation by Structure-Aware seq2seq Learning. In *Thirty-Second AAAI Conference on Artificial Intelligence*.
- [194] Tianyu Liu, Xin Zheng, Baobao Chang, and Zhifang Sui. 2021b. Towards Faithfulness in Open Domain Table-to-Text Generation from an Entity-Centric View. arXiv:2102.08585.
- [195] Yang Liu and Mirella Lapata. 2018. Learning Structured Text Representations. *Transactions of the Association for Computational Linguistics* 6 (2018), 63–75.
- [196] Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019c. Roberta: A Robustly Optimized BERT Pretraining Approach. arXiv:1907.11692.
- [197] Robert Logan, Nelson F. Liu, Matthew E. Peters, Matt Gardner, and Sameer Singh. 2019. Barack’s Wife Hillary: Using Knowledge Graphs for Fact-Aware Language Modeling. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*. 5962–5971.
- [198] Ryan Lowe, Nissan Pow, Iulian Vlad Serban, and Joelle Pineau. 2015. The Ubuntu Dialogue Corpus: A Large Dataset for Research in Unstructured Multi-Turn Dialogue Systems. In *Proceedings of the 16th Annual Meeting of the Special Interest Group on Discourse and Dialogue*. 285–294.
- [199] Yi Luan, Luheng He, Mari Ostendorf, and Hannaneh Hajishirzi. 2018. Multi-Task Identification of Entities, Relations, and Coreference for Scientific Knowledge Graph Construction. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*. 3219–3232.
- [200] Jelena Luketina, Nantas Nardelli, Gregory Farquhar, Jakob N Foerster, Jacob Andreas, Edward Grefenstette, Shimon Whiteson, and Tim Rocktäschel. 2019. A Survey of Reinforcement Learning Informed by Natural Language. In *Proceedings of the 28th International joint conference on artificial intelligence (IJCAI)*.
- [201] Minh-Thang Luong, Hieu Pham, and Christopher D Manning. 2015. Effective Approaches to Attention-based Neural Machine Translation. In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*. 1412–1421.
- [202] Shuming Ma, Pengcheng Yang, Tianyu Liu, Peng Li, Jie Zhou, and Xu Sun. 2019. Key Fact as Pivot: A Two-Stage Model for Low Resource Table-to-Text Generation. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*. 2047–2057.
- [203] Manuel Mager, Ramón Fernández Astudillo, Tahira Naseem, Md Arafat Sultan, Young-Suk Lee, Radu Florian, and Salim Roukos. 2020. GPT-too: A Language-Model-First Approach for AMR-to-Text Generation. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*. 1846–1852.
- [204] François Mairesse, Milica Gasic, Filip Jurcicek, Simon Keizer, Blaise Thomson, Kai Yu, and Steve Young. 2010. Phrase-Based Statistical Language Generation using Graphical Models and Active Learning. In *Proceedings of the 48th Annual Meeting of the Association for Computational Linguistics*. 1552–1561.
- [205] Eric Malmi, Sebastian Krause, Sascha Rothe, Daniil Mirylenka, and Aliaksei Severyn. 2019. Encode, Tag, Realize: High-Precision Text Editing. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*. 5054–5065.
- [206] Junhua Mao, Jonathan Huang, Alexander Toshev, Oana Camburu, Alan L. Yuille, and Kevin Murphy. 2016. Generation and Comprehension of Unambiguous Object Descriptions. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 11–20.

- [207] Diego Marcheggiani and Laura Perez-Beltrachini. 2018. Deep Graph Convolutional Encoders for Structured Data to Text Generation. In *Proceedings of the 11th International Conference on Natural Language Generation*. 1–9.
- [208] Diego Marcheggiani and Ivan Titov. 2017. Encoding Sentences with Graph Convolutional Networks for Semantic Role Labeling. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*. 1506–1515.
- [209] Nitika Mathur, Timothy Baldwin, and Trevor Cohn. 2019. Putting Evaluation in Context: Contextual Embeddings Improve Machine Translation Evaluation. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*. 2799–2808.
- [210] Hongyuan Mei, Mohit Bansal, and Matthew R. Walter. 2016. What to talk about and how? Selective Generation using LSTMs with Coarse-to-Fine Alignment. In *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*. 720–730.
- [211] Pablo N. Mendes, Max Jakob, and Christian Bizer. 2012. *DBpedia: A Multilingual Cross-Domain Knowledge Base*. European Language Resources Association (ELRA).
- [212] Stephen Merity, Caiming Xiong, James Bradbury, and Richard Socher. 2016. Pointer Sentinel Mixture Models. arXiv:1609.07843.
- [213] Simon Mille, Anja Belz, Bernd Bohnet, Yvette Graham, Emily Pitler, and Leo Wanner. 2018. The First Multilingual Surface Realisation Shared Task (SR'18): Overview and Evaluation Results. In *Proceedings of the First Workshop on Multilingual Surface Realisation*. 1–12.
- [214] Margaret Mitchell, Simone Wu, Andrew Zaldivar, Parker Barnes, Lucy Vasserman, Ben Hutchinson, Elena Spitzer, Inioluwa Deborah Raji, and Timnit Gebru. 2019. Model Cards for Model Reporting. In *Proceedings of the Conference on Fairness, Accountability, and Transparency*. 220–229.
- [215] Will Monroe, Robert X. D. Hawkins, Noah D. Goodman, and Christopher Potts. 2017. Colors in Context: A Pragmatic Neural Model for Grounded Language Understanding. *Transactions of the Association for Computational Linguistics* 5 (2017), 325–338.
- [216] Seungwhan Moon, Pararth Shah, Anuj Kumar, and Rajen Subba. 2019. Opendialkg: Explainable Conversational Reasoning with Attention-Based Walks over Knowledge Graphs. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*. 845–854.
- [217] Raymond J. Mooney. 2007. Learning for Semantic Parsing. In *International Conference on Intelligent Text Processing and Computational Linguistics*. Springer, 311–324.
- [218] Nafise Sadat Moosavi, Andreas Rücklé, Dan Roth, and Iryna Gurevych. 2021. Learning to Reason for Text Generation from Scientific Tables. arXiv:2104.08296.
- [219] Amit Moryossef, Yoav Goldberg, and Ido Dagan. 2019. Step-by-Step: Separating Planning from Realization in Neural Data-to-Text Generation. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, Vol. 1, Long and Short Papers. 2267–2277.
- [220] Soichiro Murakami, Akihiko Watanabe, Akira Miyazawa, Keiichi Goshima, Toshihiko Yanase, Hiroya Takamura, and Yusuke Miyao. 2017. Learning to Generate Market Comments from Stock Prices. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics*, Vol. 1, Long Papers. Association for Computational Linguistics, 1374–1384. DOI: <https://doi.org/10.18653/v1/P17-1126>
- [221] Kevin P. Murphy. 2002. Hidden semi-markov models (HSMMs). Technical report, Massachusetts Institute of Technology.
- [222] Kevin P. Murphy. 2012. *Machine Learning: A Probabilistic Perspective*. MIT press.
- [223] Linyong Nan, Lorenzo Jaime Flores, Yilun Zhao, Yixin Liu, Luke Benson, Weijin Zou, and Dragomir Radev. 2022. R2D2: Robust Data-to-Text with Replacement Detection. In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*. Association for Computational Linguistics, Abu Dhabi, United Arab Emirates, 6903–6917. Retrieved from <https://aclanthology.org/2022.emnlp-main.464>
- [224] Linyong Nan, Dragomir Radev, Rui Zhang, Amrit Rau, Abhinand Sivaprasad, Chiachun Hsieh, Xiangru Tang, Aadit Vyas, Neha Verma, Pranav Krishna, Yangxiaokang Liu, Nadia Irwanto, Jessica Pan, Faiaz Rahman, Ahmad Zaidi, Mutethia Mutuma, Yasin Tarabar, Ankit Gupta, Tao Yu, Yi Chern Tan, Xi Victoria Lin, Caiming Xiong, Richard Socher, and Nazneen Fatema Rajani. 2021. DART: Open-Domain Structured Data Record to Text Generation. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*. 432–447.
- [225] Courtney Napoles, Matthew R. Gormley, and Benjamin Van Durme. 2012. Annotated Gigaword. In *Proceedings of the Joint Workshop on Automatic Knowledge Base Construction and Web-scale Knowledge Extraction (AKBC-WEKEX)*. 95–100.
- [226] Shashi Narayan and Claire Gardent. 2012. Structure-Driven Lexicalist Generation. In *Proceedings of COLING 2012*. 2027–2042.
- [227] Neha Nayak, Dilek Hakkani-Tür, Marilyn A. Walker, and Larry P. Heck. 2017. To Plan or not to Plan? Discourse Planning in Slot-Value Informed Sequence to Sequence Models for Language Generation. In *Conference of the International Speech Communication Association*. 3339–3343.

- [228] Ani Nenkova and Kathleen McKeown. 2011. *Automatic Summarization*. Now Publishers Inc.
- [229] Feng Nie, Jinpeng Wang, Jin-ge Yao, Rong Pan, and Chin-Yew Lin. 2018. Operation-Guided Neural Networks for High Fidelity Data-To-Text Generation. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*. 3879–3889.
- [230] Jekaterina Novikova, Ondrej Dusek, and Verena Rieser. 2017. The E2E Dataset: New Challenges For End-to-End Generation. In *Proceedings of the 18th Annual Meeting of the Special Interest Group on Discourse and Dialogue*. Association for Computational Linguistics, 201–206.
- [231] Jekaterina Novikova, Ondřej Dušek, and Verena Rieser. 2018. RankME: Reliable Human Ratings for Natural Language Generation. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, Vol. 2, Short Papers. 72–78.
- [232] Jekaterina Novikova, Oliver Lemon, and Verena Rieser. 2016. Crowd-Sourcing NLG Data: Pictures Elicit Better Data. In *Proceedings of the 9th International Natural Language Generation Conference*. 265–273.
- [233] Jason Obeid and Enamul Hoque. 2020. Chart-to-Text: Generating Natural Language Descriptions for Charts by Adapting the Transformer Model. In *Proceedings of the 13th International Conference on Natural Language Generation*. 138–147.
- [234] Hugo Gonçalo Oliveira. 2017. A Survey on Intelligent Poetry Generation: Languages, Features, Techniques, Re-utilisation and Evaluation. In *Proceedings of the 10th International Conference on Natural Language Generation*. 11–20.
- [235] OpenAI. 2022. Chat-GPT: Optimizing Language Models for Dialogue. Retrieved from <https://openai.com/blog/chatgpt/>
- [236] Shereen Oraby, Vrindavan Harrison, Abteen Ebrahimi, and Marilyn Walker. 2019. Curate and Generate: A Corpus and Method for Joint Control of Semantics and Style in Neural NLG. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*.
- [237] Shereen Oraby, Lena Reed, Shubhangi Tandon, TS Sharath, Stephanie Lukin, and Marilyn Walker. 2018. Controlling Personality-Based Stylistic Variation with Neural Natural Language Generators. In *Proceedings of the 19th Annual SIGdial Meeting on Discourse and Dialogue*. 180–190.
- [238] Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. Bleu: A Method for Automatic Evaluation of Machine Translation. In *Proceedings of the 40th Annual Meeting of the Association for Computational Linguistics*. 311–318.
- [239] Ankur Parikh, Xuezhi Wang, Sebastian Gehrmann, Manaal Faruqui, Bhuwan Dhingra, Diyi Yang, and Dipanjan Das. 2020. ToTTo: A Controlled Table-To-Text Generation Dataset. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*. 1173–1186.
- [240] Romain Paulus, Caiming Xiong, and Richard Socher. 2018. A Deep Reinforced Model for Abstractive Summarization. In *International Conference on Learning Representations*.
- [241] Steffen Pauws, Albert Gatt, Emiel Krahmer, and Ehud Reiter. 2019. Making Effective Use of Healthcare Data using Data-to-Text Technology. In *Data Science for Healthcare*. Springer, 119–145.
- [242] Judea Pearl. 2010. On Measurement Bias in Causal Inference. In *Proceedings of the Twenty-Sixth Conference on Uncertainty in Artificial Intelligence (UAI'10)*. AUAI Press, Arlington, VA, USA, 425–432.
- [243] Jaclyn Peiser. 2019. The Rise of the Robot Reporter. *The New York Times* (Feb 2019).
- [244] Laura Perez-Beltrachini and Mirella Lapata. 2018. Bootstrapping Generators from Noisy Data. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, Vol. 1, Long Papers. 1516–1527.
- [245] Matthew E. Peters, Mark Neumann, Mohit Iyyer, Matt Gardner, Christopher Clark, Kenton Lee, and Luke Zettlemoyer. 2018. Deep Contextualized Word Representations. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, Vol. 1, Long Papers. Association for Computational Linguistics, 2227–2237. DOI: <https://doi.org/10.18653/v1/N18-1202>
- [246] Marian Petre. 1995. Why Looking isn't Always Seeing: Readership Skills and Graphical Programming. *Communications of the ACM* 38, 6 (1995), 33–44.
- [247] Fabio Petroni, Tim Rocktäschel, Sebastian Riedel, Patrick Lewis, Anton Bakhtin, Yuxiang Wu, and Alexander Miller. 2019. Language Models as Knowledge Bases?. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*. 2463–2473.
- [248] Joelle Pineau, Philippe Vincent-Lamarre, Koustuv Sinha, Vincent Larivière, Alina Beygelzimer, Florence d'Alché Buc, Emily Fox, and Hugo Larochelle. 2021. Improving Reproducibility in Machine Learning Research (a Report from the NeurIPS 2019 Reproducibility Program). *The Journal of Machine Learning Research* 22, 1 (2021), 7459–7478.
- [249] François Portet, Ehud Reiter, Albert Gatt, Jim Hunter, Somayajulu Sripada, Yvonne Freer, and Cindy Sykes. 2009. Automatic Generation of Textual Summaries from Neonatal Intensive Care Data. *Artificial Intelligence* 173, 7-8 (2009), 789–816.



- [250] Matt Post. 2018. A Call for Clarity in Reporting BLEU Scores. In *Proceedings of the Third Conference on Machine Translation: Research Papers*. 186–191.
- [251] Ellen F. Prince. 1981. Towards a taxonomy of given-new information. In *Radical Pragmatics*. P. Coli (Ed.), Academic Press, New York, NY, USA, 223–256.
- [252] Ratish Puduppully, Li Dong, and Mirella Lapata. 2019a. Data-to-Text Generation with Content Selection and Planning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 33. 6908–6915.
- [253] Ratish Puduppully, Li Dong, and Mirella Lapata. 2019b. Data-to-text Generation with Entity Modeling. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*. 2023–2035.
- [254] Ratish Puduppully, Yao Fu, and Mirella Lapata. 2022. Data-to-text Generation with Variational Sequential Planning. arXiv:2202.13756.
- [255] Ratish Puduppully and Mirella Lapata. 2021. Data-to-text Generation with Macro Planning. *Transactions of the Association for Computational Linguistics* 9 (2021), 510–527.
- [256] Yevgeniy Puzikov and Iryna Gurevych. 2018. E2e nlg Challenge: Neural Models vs. Templates. In *Proceedings of the 11th International Conference on Natural Language Generation*. 463–471.
- [257] Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. 2019. Language Models are Unsupervised Multitask Learners. *OpenAI Blog* 1, 8 (2019), 9.
- [258] Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J. Liu. 2019. Exploring the Limits of Transfer Learning with a Unified Text-to-Text Transformer. arXiv:1910.10683.
- [259] Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J Liu. 2020. Exploring the Limits of Transfer Learning with a Unified Text-to-Text Transformer. *Journal of Machine Learning Research* 21 (2020), 1–67.
- [260] Juan Ramos. 2003. Using TF-IDF to determine Word Relevance in Document Queries. In *Proceedings of the First Instructional Conference on Machine Learning*, Vol. 242. Citeseer, 29–48.
- [261] Clément Rebuffel, Laure Soulier, Geoffrey Scuttheeten, and Patrick Gallinari. 2020a. A Hierarchical Model for Data-to-Text Generation. In *European Conference on Information Retrieval*. Springer, 65–80.
- [262] Clément Rebuffel, Laure Soulier, Geoffrey Scuttheeten, and Patrick Gallinari. 2020b. Parenting via Model-Agnostic Reinforcement Learning to correct Pathological Behaviors in Data-to-Text Generation. arXiv:2010.10866.
- [263] Lena Reed, Shereen Oraby, and Marilyn Walker. 2018. Can Neural Generators for Dialogue Learn Sentence Planning and Discourse Structuring? In *Proceedings of the 11th International Conference on Natural Language Generation*. 284–295.
- [264] Marek Rei. 2017. Semi-Supervised Multitask Learning for Sequence Labeling. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics*, Vol. 1, Long Papers. 2121–2130.
- [265] Nils Reimers and Iryna Gurevych. 2017. Reporting Score Distributions Makes a Difference: Performance Study of LSTM-Networks for Sequence Tagging. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*. 338–348.
- [266] Ehud Reiter. 2007. An Architecture for Data-to-Text Systems. In *Proceedings of the 11th European Workshop on Natural Language Generation (ENLG'07)*. 97–104.
- [267] Ehud Reiter. 2017. You Need to Understand Your Corpora! The Weathervgov example. Retrieved from <https://ehudreiter.com/2017/05/09/weathervgov/>
- [268] Ehud Reiter. 2018. A Structured Review of the Validity of BLEU. *Computational Linguistics* 44, 3 (2018), 393–401.
- [269] Ehud Reiter and Anja Belz. 2009. An Investigation into the Validity of Some Metrics for Automatically Evaluating Natural Language Generation Systems. *Computational Linguistics* 35, 4 (2009), 529–558.
- [270] Ehud Reiter and Robert Dale. 1997. Building Applied Natural Language Generation Systems. *Natural Language Engineering* 3, 1 (1997), 57–87.
- [271] Ehud Reiter, Somayajulu Sripada, Jim Hunter, Jin Yu, and Ian Davy. 2005. Choosing Words in Computer-Generated Weather Forecasts. *Artificial Intelligence* 167, 1-2 (2005), 137–169.
- [272] Steven J. Rennie, Etienne Marcheret, Youssef Mroueh, Jerret Ross, and Vaibhava Goel. 2017. Self-Critical Sequence Training for Image Captioning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 7008–7024.
- [273] Leonardo F. R. Ribeiro, Claire Gardent, and Iryna Gurevych. 2019. Enhancing AMR-to-Text Generation with Dual Graph Representations. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*. 3183–3194.
- [274] Leonardo F. R. Ribeiro, Martin Schmitt, Hinrich Schütze, and Iryna Gurevych. 2021. Investigating Pretrained Language Models for Graph-to-Text Generation. In *Proceedings of the 3rd Workshop on Natural Language Processing for Conversational AI*. 211–227.

- [275] Leonardo F. R. Ribeiro, Yue Zhang, Claire Gardent, and Iryna Gurevych. 2020. Modeling Global and Local Node Contexts for Text Generation from Knowledge Graphs. *Transactions of the Association for Computational Linguistics* 8 (2020), 589–604.
- [276] Markus Ring, Daniel Schlör, Dieter Landes, and Andreas Hotho. 2019. Flow-Based Network Traffic Generation using Generative Adversarial Networks. *Computers & Security* 82 (2019), 156–172.
- [277] Marco Roberti, Giovanni Bonetta, Rossella Cancelliere, and Patrick Gallinari. 2019. Copy Mechanism and Tailored Training for Character-Based Data-to-Text Generation. In *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*. Springer, 648–664.
- [278] Stephen Robertson and Hugo Zaragoza. 2009. The Probabilistic Relevance Framework: BM25 and Beyond. *Foundations and Trends in Information Retrieval* 3, 4 (2009), 333–389.
- [279] Jacques Pierre Robin. 1995. *Revision-Based Generation of Natural Language Summaries Providing Historical Background: Corpus-Based Analysis, Design, Implementation and Evaluation*. Ph. D. Dissertation. Columbia University.
- [280] Anna Rohrbach, Lisa Anne Hendricks, Kaylee Burns, Trevor Darrell, and Kate Saenko. 2018. Object Hallucination in Image Captioning. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*. 4035–4045.
- [281] Andreas Rücklé, Steffen Eger, Maxime Peyrard, and Iryna Gurevych. 2018. Concatenated Power Mean Word Embeddings as Universal Cross-Lingual Sentence Representations. arXiv:1803.01400.
- [282] Sashank Santhanam and Samira Shaikh. 2019. A Survey of Natural Language Generation Techniques with a Focus on Dialogue Systems-Past, Present and Future Directions. arXiv:1906.00500.
- [283] Jost Schatzmann, Blaise Thomson, Karl Weilhammer, Hui Ye, and Steve Young. 2007. Agenda-Based User Simulation for Bootstrapping a POMDP Dialogue System. In *Human Language Technologies 2007: The Conference of the North American Chapter of the Association for Computational Linguistics*, Companion Volume, Short Papers. 149–152.
- [284] Timo Schick, Jane Dwivedi-Yu, Roberto Dessò, Roberta Raileanu, Maria Lomeli, Luke Zettlemoyer, Nicola Cancedda, and Thomas Scialom. 2023. Toolformer: Language Models Can Teach Themselves to Use Tools. arXiv abs/2302.04761.
- [285] Michael Schlichtkrull, Thomas N. Kipf, Peter Bloem, Rianne van den Berg, Ivan Titov, and Max Welling. 2018. Modeling Relational Data with Graph Convolutional Networks. In *European Semantic Web Conference*. Springer, 593–607.
- [286] Donia Scott and Johanna Moore. 2007. An NLG Evaluation Competition? Eight Reasons to be Cautious. In *Proceedings of the NSF Workshop on Shared Tasks and Comparative Evaluation in Natural Language Generation*.
- [287] Abigail See, Peter J. Liu, and Christopher D. Manning. 2017. Get To The Point: Summarization with Pointer-Generator Networks. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics*, Vol. 1, Long Papers. Association for Computational Linguistics, 1073–1083. DOI: <https://doi.org/10.18653/v1/P17-1099>
- [288] Abigail See, Aneesh Pappu, Rohun Saxena, Akhila Yerukola, and Christopher D. Manning. 2019. Do Massively Pretrained Language Models Make Better Storytellers?. In *Proceedings of the 23rd Conference on Computational Natural Language Learning (CoNLL)*. 843–861.
- [289] Edward Segel and Jeffrey Heer. 2010. Narrative Visualization: Telling Stories with Data. *IEEE Transactions on Visualization and Computer Graphics* 16, 6 (2010), 1139–1148.
- [290] Rico Sennrich, Barry Haddow, and Alexandra Birch. 2016. Improving Neural Machine Translation Models with Monolingual Data. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics*, Vol. 1, Long Papers. 86–96.
- [291] Iulian Serban, Alessandro Sordani, Ryan Lowe, Laurent Charlin, Joelle Pineau, Aaron Courville, and Yoshua Bengio. 2017. A Hierarchical Latent Variable Encoder-Decoder Model for Generating Dialogues. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 31.
- [292] Lei Sha, Lili Mou, Tianyu Liu, Pascal Poupart, Sujian Li, Baobao Chang, and Zhifang Sui. 2018. Order-Planning Neural Text Generation from Structured Data. In *Proceedings of the 32nd AAAI Conference on Artificial Intelligence*.
- [293] Zhihong Shao, Minlie Huang, Jiangtao Wen, Wenfei Xu, and Xiaoyan Zhu. 2019. Long and Diverse Text Generation with Planning-based Hierarchical Variational Model. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*. 3257–3268.
- [294] Mandar Sharma, John S Brownstein, and Naren Ramakrishnan. 2021. T 3: Domain-Agnostic Neural Time-series Narration. In *IEEE International Conference on Data Mining (ICDM'21)*. IEEE, 1324–1329.
- [295] Mandar Sharma, Nikhil Muralidhar, and Naren Ramakrishnan. 2022. Overcoming Barriers to Skill Injection in Language Modeling: Case Study in Arithmetic. arXiv:2211.02098.
- [296] Peter Shaw, Jakob Uszkoreit, and Ashish Vaswani. 2018. Self-Attention with Relative Position Representations. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, Vol. 2, Short Papers. 464–468.

- [297] Sheng Shen, Daniel Fried, Jacob Andreas, and Dan Klein. 2019. Pragmatically Informative Text Generation. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, Vol. 1, Long and Short Papers. 4060–4067.
- [298] Stuart Shieber, Gertjan Van Noord, Fernando C. N. Pereira, and Robert C. Moore. 1990. Semantic-Head-Driven Generation. *Computational Linguistics* 16, 1 (1990), 30–41.
- [299] Anastasia Shimorina and Anja Belz. 2022. The Human Evaluation Datasheet: A Template for Recording Details of Human Evaluation Experiments in NLP. In *Proceedings of the Second Workshop on Human Evaluation of NLP Systems (HumEval)*. 54–75.
- [300] Anastasia Shimorina and Claire Gardent. 2018. Handling Rare Items in Data-to-Text Generation. In *Proceedings of the 11th International Conference on Natural Language Generation*. 360–370.
- [301] Advait Siddharthan and Napoleon Katsos. 2012. Offline Sentence Processing Measures for Testing Readability with Users. In *Proceedings of the First Workshop on Predicting and Improving Text Readability for Target Reader Populations*. 17–24.
- [302] Koustuv Sinha, Joelle Pineau, Jessica Forde, Rosemary Nan Ke, and Hugo Larochelle. 2020. NeurIPS 2019 Reproducibility Challenge. *ReScience C* 6, 2 (2020), 11.
- [303] Charese Smiley, Vassilis Plachouras, Frank Schilder, Hiroko Bretz, Jochen L. Leidner, and Dezhao Song. 2016. When to Plummet and When to Soar: Corpus Based Verb Selection for Natural Language Generation. In *Proceedings of the 9th International Natural Language Generation Conference*. 36–39.
- [304] Richard Socher, Jeffrey Pennington, Eric H. Huang, Andrew Y. Ng, and Christopher D. Manning. 2011. Semi-Supervised Recursive Autoencoders for Predicting Sentiment Distributions. In *Proceedings of the 2011 Conference on Empirical Methods in Natural Language Processing*. 151–161.
- [305] Kihyuk Sohn, Honglak Lee, and Xinchen Yan. 2015. Learning Structured Output Representation using Deep Conditional Generative Models. *Advances in Neural Information Processing Systems* 28 (2015).
- [306] Linfeng Song, Ante Wang, Jinsong Su, Yue Zhang, Kun Xu, Yubin Ge, and Dong Yu. 2020. Structural Information Preserving for Graph-to-Text Generation. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*. 7987–7998.
- [307] Amanda Stent, Matthew Marge, and Mohit Singhai. 2005. Evaluating Evaluation Methods for Generation in the Presence of Variation. In *International Conference on Intelligent Text Processing and Computational Linguistics*. Springer, 341–351.
- [308] Amanda Stent, Rashmi Prasad, and Marilyn Walker. 2004. Trainable Sentence Planning for Complex Information Presentations in Spoken Dialog Systems. In *Proceedings of the 42nd Annual Meeting of the Association for Computational Linguistics (ACL'04)*. 79–86.
- [309] Milan Straka and Jana Straková. 2017. Tokenizing, Pos Tagging, Lemmatizing and Parsing Ud 2.0 with Udpipes. In *Proceedings of the CoNLL 2017 Shared Task: Multilingual Parsing from Raw Text to Universal Dependencies*. 88–99.
- [310] Shang-Yu Su, Kai-Ling Lo, Yi-Ting Yeh, and Yun-Nung Chen. 2018. Natural Language Generation by Hierarchical Decoding with Linguistic Patterns. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, Vol. 2, Short Papers. 61–66.
- [311] Yixuan Su, Zaiqiao Meng, Simon Baker, and Nigel Collier. 2021. Few-Shot Table-to-Text Generation with Prototype Memory. In *Findings of the Association for Computational Linguistics: EMNLP'21*. 910–917.
- [312] Lya Hulliyiyatus Suadaa, Hidetaka Kamigaito, Kotaro Funakoshi, Manabu Okumura, and Hiroya Takamura. 2021. Towards Table-to-Text Generation with Numerical Reasoning. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing*, Vol. 1, Long Papers. 1451–1465.
- [313] Simeng Sun, Ori Shapira, Ido Dagan, and Ani Nenkova. 2019. How to Compare Summarizers without Target Length? Pitfalls, Solutions and Re-Examination of the Neural Summarization Literature. In *Proceedings of the Workshop on Methods for Optimizing and Evaluating Neural Language Generation*. 21–29.
- [314] Ilya Sutskever, James Martens, and Geoffrey E. Hinton. 2011. Generating text with recurrent neural networks. In *Proceedings of the 28th International Conference on Machine Learning*.
- [315] Richard S. Sutton and Andrew G. Barto. 2018. *Reinforcement Learning: An Introduction*. MIT press.
- [316] Kumiko Tanaka-Ishii, Kôiti Hasida, and Itsuki Noda. 1998. Reactive Content Selection in the Generation of Real-Time Soccer Commentary. In *COLING 1998 Volume 2: The 17th International Conference on Computational Linguistics*.
- [317] Avijit Thawani, Jay Pujara, Filip Ilievski, and Pedro Szekely. 2021. Representing Numbers in NLP: a Survey and a Vision. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*. 644–656.

- [318] Ran Tian, Shashi Narayan, Thibault Sellam, and Ankur P. Parikh. 2019. Sticking to the Facts: Confident Decoding for Faithful Data-to-Text Generation. arXiv:1910.08684d.
- [319] Jakub Tomczak and Max Welling. 2018. VAE with a VampPrior. In *International Conference on Artificial Intelligence and Statistics*. PMLR, 1214–1223.
- [320] Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, Aurelien Rodriguez, Armand Joulin, Edouard Grave, and Guillaume Lample. 2023. LLaMA: Open and Efficient Foundation Language Models. arXiv:2302.13971.
- [321] Zhaopeng Tu, Yang Liu, Lifeng Shang, Xiaohua Liu, and Hang Li. 2017. Neural Machine Translation with Reconstruction. In *Proceedings of the 31st AAAI Conference on Artificial Intelligence*.
- [322] Zhaopeng Tu, Zhengdong Lu, Yang Liu, Xiaohua Liu, and Hang Li. 2016. Modeling Coverage for Neural Machine Translation. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics*, Vol. 1, Long Papers. 76–85.
- [323] Yui Uehara, Tatsuya Ishigaki, Kasumi Aoki, Hiroshi Noji, Keiichi Goshima, Ichiro Kobayashi, Hiroya Takamura, and Yusuke Miyao. 2020. Learning with Contrastive Examples for Data-to-Text Generation. In *Proceedings of the 28th International Conference on Computational Linguistics*. 2352–2362.
- [324] Chris van der Lee, Chris Emmerly, Sander Wubben, and Emiel Kraahmer. 2020. The CACAPO Dataset: A Multilingual, Multi-Domain Dataset for Neural Pipeline and End-to-End Data-to-Text Generation. In *Proceedings of the 13th International Conference on Natural Language Generation*. 68–79.
- [325] Chris van der Lee, Albert Gatt, Emiel van Miltenburg, and Emiel Kraahmer. 2021. Human Evaluation of Automatically Generated Text: Current Trends and Best Practice Guidelines. *Computer Speech & Language* 67 (2021), 101151.
- [326] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. *Advances in Neural Information Processing Systems* 30 (2017), 15.
- [327] Petar Veličković, Guillem Cucurull, Arantxa Casanova, Adriana Romero, Pietro Lio, and Yoshua Bengio. 2018. Graph Attention Networks. In *International Conference on Learning Representations*.
- [328] Subhashini Venugopalan, Marcus Rohrbach, Jeffrey Donahue, Raymond Mooney, Trevor Darrell, and Kate Saenko. 2015. Sequence to Sequence-Video to Text. In *Proceedings of the IEEE International Conference on Computer Vision*. 4534–4542.
- [329] Pascal Vincent, Hugo Larochelle, Yoshua Bengio, and Pierre-Antoine Manzagol. 2008. Extracting and Composing Robust Features with Denoising Autoencoders. In *Proceedings of the 25th International Conference on Machine Learning*. 1096–1103.
- [330] Pascal Vincent, Hugo Larochelle, Isabelle Lajoie, Yoshua Bengio, Pierre-Antoine Manzagol, and Léon Bottou. 2010. Stacked Denoising Autoencoders: Learning Useful Representations in a Deep Network with a Local Denoising Criterion. *Journal of Machine Learning Research* 11, 12 (2010), 3371–3408.
- [331] Oriol Vinyals, Meire Fortunato, and Navdeep Jaitly. 2015. Pointer Networks. *Advances in Neural Information Processing Systems* 28 (2015).
- [332] Denny Vrandečić and Markus Krötzsch. 2014. Wikidata: A Free Collaborative Knowledgebase. *Communications of the ACM* 57, 10 (2014), 78–85.
- [333] Marilyn A. Walker, Amanda Stent, François Mairesse, and Rashmi Prasad. 2007. Individual and Domain Adaptation in Sentence Planning for Dialogue. *Journal of Artificial Intelligence Research* 30 (2007), 413–456.
- [334] Eric Wallace, Yizhong Wang, Sujian Li, Sameer Singh, and Matt Gardner. 2019. Do NLP Models Know Numbers? Probing Numeracy in Embeddings. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*. 5307–5315.
- [335] Hongmin Wang. 2019. Revisiting Challenges in Data-to-Text Generation with Fact Grounding. In *Proceedings of the 12th International Conference on Natural Language Generation*. 311–322.
- [336] Jun Wang and Klaus Mueller. 2015. The Visual Causality Analyst: An Interactive Interface for Causal Reasoning. *IEEE Transactions on Visualization and Computer Graphics* 22, 1 (2015), 230–239.
- [337] Peng Wang, Junyang Lin, An Yang, Chang Zhou, Yichang Zhang, Jingren Zhou, and Hongxia Yang. 2021. Sketch and Refine: Towards Faithful and Informative Table-to-Text Generation. In *Findings of the Association for Computational Linguistics: ACL-IJCNLP.21*. 4831–4843.
- [338] Qingyun Wang, Xiaoman Pan, Lifu Huang, Boliang Zhang, Zhiying Jiang, Heng Ji, and Kevin Knight. 2018. Describing a Knowledge Base. In *Proceedings of the 11th International Conference on Natural Language Generation*. 10–21.
- [339] Zhenyi Wang, Xiaoyang Wang, Bang An, Dong Yu, and Changyou Chen. 2020. Towards Faithful Neural Table-to-Text Generation with Content-Matching Constraints. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*. 1072–1086.
- [340] Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Fei Xia, Ed H. Chi, Quoc V. Le, and Denny Zhou. 2022. Chain-of-Thought Prompting Elicits Reasoning in Large Language Models. In *Advances in Neural Information Processing Systems*.

- [341] Sean Welleck, Ilya Kulikov, Stephen Roller, Emily Dinan, Kyunghyun Cho, and Jason Weston. 2019. Neural Text Generation With Unlikelihood Training. In *International Conference on Learning Representations*.
- [342] Tsung-Hsien Wen, Milica Gasic, Nikola Mrkšić, Lina Maria Rojas-Barahona, Pei-Hao Su, David Vandyke, and Steve J. Young. 2016. Multi-Domain Neural Network Language Generation for Spoken Dialogue Systems. In *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*.
- [343] Tsung-Hsien Wen, Milica Gasic, Nikola Mrkšić, Pei-Hao Su, David Vandyke, and Steve Young. 2015. Semantically Conditioned LSTM-based Natural Language Generation for Spoken Dialogue Systems. In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, 1711–1721.
- [344] Lilian Weng. 2021. Controllable Neural Text Generation. *lilianweng.github.io*. Retrieved from <https://lilianweng.github.io/posts/2021-01-02-controllable-text-generation/>
- [345] Ronald J. Williams and Jing Peng. 1991. Function Optimization using Connectionist Reinforcement Learning Algorithms. *Connection Science* 3, 3 (1991), 241–268.
- [346] Ronald J. Williams and David Zipser. 1989. A Learning Algorithm for Continually Running Fully Recurrent Neural Networks. *Neural cMputation* 1, 2 (1989), 270–280.
- [347] Sam Wiseman, Stuart M. Shieber, and Alexander M. Rush. 2017. Challenges in Data-to-Document Generation. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*.
- [348] Sam Wiseman, Stuart M. Shieber, and Alexander M. Rush. 2018. Learning Neural Templates for Text Generation. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, 3174–3187.
- [349] Stephen Wolfram. 2023. Wolfram—Alpha as the Way to Bring Computational Knowledge Superpowers to ChatGPT. Retrieved from <https://writings.stephenwolfram.com/2023/01/wolframalpha-as-the-way-to-bring-computational-knowledge-superpowers-to-chatgpt/>
- [350] Yonghui Wu, Mike Schuster, Zhifeng Chen, Quoc V Le, Mohammad Norouzi, Wolfgang Macherey, Maxim Krikun, Yuan Cao, Qin Gao, Klaus Macherey, Jeff Klingner, Apurva Shah, Melvin Johnson, Xiaobing Liu, Łukasz Kaiser, Stephan Gouws, Yoshikiyo Kato, Taku Kudo, Hideto Kazawa, Keith Stevens, George Kurian, Nishant Patil, Wei Wang, Cliff Young, Jason Smith, Jason Riesa, Alex Rudnick, Oriol Vinyals, Greg Corrado, Macduff Hughes, and Jeffrey Dean. 2016. Google’s Neural Machine Translation System: Bridging the Gap between Human and Machine Translation. Retrieved from <https://doi.org/10.48550/arXiv.1609.08144>
- [351] Jiannan Xiang, Zhengzhong Liu, Yucheng Zhou, Eric Xing, and Zhiting Hu. 2022. ASDOT: Any-Shot Data-to-Text Generation with Pretrained Language Models. In *Findings of the Association for Computational Linguistics: EMNLP 2022*. Association for Computational Linguistics, 1886–1899.
- [352] Kelvin Xu, Jimmy Ba, Ryan Kiros, Kyunghyun Cho, Aaron Courville, Ruslan Salakhudinov, Rich Zemel, and Yoshua Bengio. 2015. Show, Attend and Tell: Neural Image Caption Generation with Visual Attention. In *International Conference on Machine Learning*. PMLR, 2048–2057.
- [353] Shengzhe Xu, Manish Marwah, Martin Arlitt, and Naren Ramakrishnan. 2021. Stan: Synthetic Network Traffic Generation with Generative Neural Models. In *Deployable Machine Learning for Security Defense: Second International Workshop, MLHat 2021, Virtual Event, Proceedings 2*. Springer, 3–29.
- [354] Yang Yang, Juan Cao, Yujun Wen, and Pengzhou Zhang. 2021. Table to Text Generation with Accurate Content Copying. *Scientific Reports* 11, 1 (2021), 1–12.
- [355] Zichao Yang, Phil Blunsom, Chris Dyer, and Wang Ling. 2017. Reference-Aware Language Models. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, 1850–1859.
- [356] Rong Ye, Wenxian Shi, Hao Zhou, Zhongyu Wei, and Lei Li. 2020. Variational Template Machine for Data-to-Text Generation. Retrieved from <https://doi.org/10.48550/arXiv.2002.01127>
- [357] Steve Young, Milica Gašić, Simon Keizer, François Mairesse, Jost Schatzmann, Blaise Thomson, and Kai Yu. 2010. The Hidden Information State Model: A Practical Framework for POMDP-Based Spoken Dialogue Management. *Computer Speech & Language* 24, 2 (2010), 150–174.
- [358] Wlodek Zadrozny and Karen Jensen. 1991. Semantics of Paragraphs. *Computational Linguistics* 17, 2 (1991), 171–210.
- [359] Biao Zhang, Deyi Xiong, Jinsong Su, and Hong Duan. 2017. A Context-Aware Recurrent Encoder for Neural Machine Translation. *IEEE/ACM Transactions on Audio, Speech, and Language Processing* 25, 12 (2017), 2424–2432.
- [360] Biao Zhang, Jing Yang, Qian Lin, and Jinsong Su. 2018. Attention Regularized Sequence-to-Sequence Learning for E2E NLG Challenge. In *Proceedings of the E2E NLG Challenge System Descriptions*.
- [361] Shuo Zhang, Zhuyun Dai, Krisztian Balog, and Jamie Callan. 2020a. Summarizing and Exploring Tabular Data in Conversational Search. In *Proceedings of the 43rd International ACM SIGIR Conference on Research and Development in Information Retrieval*, 1537–1540.
- [362] Tianyi Zhang, Varsha Kishore, Felix Wu, Kilian Q. Weinberger, and Yoav Artzi. 2019. BERTScore: Evaluating Text Generation with BERT. In *International Conference on Learning Representations*.

- [363] Yizhe Zhang, Siqi Sun, Michel Galley, Yen-Chun Chen, Chris Brockett, Xiang Gao, Jianfeng Gao, Jingjing Liu, and William B. Dolan. 2020b. DIALOGPT: Large-Scale Generative Pre-training for Conversational Response Generation. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics: System Demonstrations*. 270–278.
- [364] Chao Zhao, Marilyn Walker, and Snigdha Chaturvedi. 2020. Bridging the Structural Gap Between Encoding and Decoding for Data-to-Text Generation. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*. 2481–2491.
- [365] Jianyu Zhao, Zhiqiang Zhan, Tong Li, Rang Li, Changjian Hu, Siyun Wang, and Yang Zhang. 2021. Generative Adversarial Network for Table-to-Text Generation. *Neurocomputing* 452 (2021), 28–36.
- [366] Wei Zhao, Maxime Peyrard, Fei Liu, Yang Gao, Christian M Meyer, and Steffen Eger. 2019. MoverScore: Text Generation Evaluating with Contextualized Embeddings and Earth Mover Distance. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*. 563–578.
- [367] Jie Zhu, Junhui Li, Muhua Zhu, Longhua Qian, Min Zhang, and Guodong Zhou. 2019. Modeling Graph Structure in Transformer for Better AMR-to-Text Generation. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*. 5459–5468.
- [368] Brian D. Ziebart, Andrew L. Maas, J. Andrew Bagnell, and Anind K. Dey. 2008. Maximum Entropy Inverse Reinforcement Learning. In *AAAI*, Vol. 8. 1433–1438.

Received 12 April 2023; revised 26 March 2024; accepted 28 March 2024