A Supplementary Material

show, tell, did, me, my, our, are, is, were, this, on,
would, and, for, should, be, do, I, have, had, the,
there, look, give, has, was, we, get, does, a, an, 's,
that, by, based, in, of, bring, with, to, from, whole,
being, been, want, wanted, as, can, see, doing,
got, sorted, draw, listed, chart, only

Figure 5: Neural Programmer (Neelakantan et al., 2017): List of stop words used in stop word deletion attacks (section 5.5). NP's accuracy falls from 33.5% to 28.5% on deleting stop words from questions in the validation set.

Operator	Triggers
select	[tm_token, many, how, number, or, total, after, before, only]
prev	[before, many, than, previous, above, how, at, most]
first	[tm_token, first, before, after, who, previous, or, peak]
reset	[many, total, how, number, last, least, the, first, of]
count	[many, how, number, total, of, difference, between, long, times]
next	[after, not, many, next, same, tm_token, how, below]
last	[last, or, after, tm_token, next, the, chart, not]
mfe	[most, cm_token, same]
min	[least, the, not]
max	[most, largest]
geq	[at, more, least, had, over, number, than, many]
print	[tm_token]

Table 5: Neural Programmer (Neelakantan et al., 2017): Operator triggers. Notice that there are several irrelevant triggers (highlighted in red). For instance, "many" is irrelevant to "prev". See Section 5.5 for attacks exploiting this weakness.

Operator sequence	#tables	Triggers
reset, reset, max, print	109	[*unk*, date, position, points, name, competition, notes, no, year, venue]
reset, prev, max, print	68	[*unk*, rank, total, bronze, gold, silver, nation, name, date, no]
reset, reset, first, print	29	[name, *unk*, notes, year, nationality, rank, location, date, comments, hometown]
reset, mfe, first, print	25	[notes, date, title, *unk*, role, genre, year, score, opponent, event]
reset, reset, min, print	17	[year, height, *unk*, name, position, floors, notes, jan, jun, may]
reset, mfe, max, print	14	[opponent, date, result, location, rank, site, attendance, notes, city, listing]
reset, next, first, print	10	[*unk*, name, year, edition, birth, death, men, time, women, type]
reset, reset, last, print	9	[date, *unk*, distance, location, name, year, winner, japanese, duration, member]
reset, prev, first, print	7	[name, notes, intersecting, kilometers, location, athlete, nationality, rank, time, design]
reset, next, max, print	7	[*unk*, ethnicity]
reset, mfe, last, print	6	[place, season, *unk*, date, division, tier, builder, cylinders, notes, withdrawn]
reset, prev, last, print	5	[report, date, average, chassis, driver, race, builder, name, notes, works]
reset, prev, min, print	4	[division, level, movements, position, season, current, gauge, notes, wheel, works]
reset, select, last, print	3	[car, finish, laps, led, rank, retired, start, year, *unk*, carries]
reset, reset, first, count	2	[*unk*, network, owner, programming]
reset, reset, reset, print	1	[candidates, district, first, incumbent, party, result]
reset, select, select, print	1	[lifetime, name, nationality, notable, notes]
reset, reset, last, count	1	[*unk*, english, japanese, type, year]
reset, next, last, print	1	[*unk*, comment]
reset, mfe, select, print	1	[*unk*, length, performer, producer, title]

Table 6: Neural Programmer (Neelakantan et al., 2017): Column attributions in occurring in table-specific default programs, classified by operator sequence. These attributions are indication that the network is predisposed towards picking certain operators solely based on the table. Here is an insight: the second row of this table indicates that the network is inclined to choosing "reset, prev" on tables about medal tallies. It may have learned this bias as some medal tables in training data may have "*total*" rows which may confound answer computation if not excluded by applying "reset, prev". However, not all medal tables have "*total*" rows, and hence "reset, prev" must not be applied universally. As the operators are triggered by column names and not by table entries, the network cannot distinguish between tables with and without "*total*" row and may erroneously exclude the last rows.