Bachelor Thesis Review

Faculty of Mathematics and Physics, Charles University

Thesis author	Sviatoslav Gladkykh
Thesis title	Denoising Diffusion Models for Dynamic Sky Image Generation
Year submitted	2024
Study program	Computer Science
Specialization	Artificial Intelligence
Review author	Mgr. Martin Mirbauer — Reviewer
Department	Department of Software and Computer Science Education

Overall

	good	OK	poor	insufficient
Assignment difficulty		X		
Assignment fulfilled	X	X		
Total size *... text and code, overall workload*		X		

The author successfully applied diffusion models to the sky generation task and achieved good results without significant artefacts. I noticed only slight deformation of the sun/flares shape and a tendency of generating slightly more blurry/grey images compared to the training images, when generating the 20-frames-long sequences.

Thesis Text

	good	OK	poor	insufficient
Form *... language, typography, references*		X	X	
Structure *... context, goals, analysis, design, evaluation, level of detail*		X		
Problem analysis		X		
Developer documentation			X	
User Documentation			X	

A more detailed description of the results would be welcome. E.g. a side-by-side comparison with nearest neighbour from the training dataset to rule out overfitting, or a discussion of possible artefacts.

I would expect a clearer description of which parts of the architecture and also code are pre-existing (e.g. diffusion *backbone*), and which were adapted or written from scratch.

The comparison of FID to previous work (section 8) may not be relevant, since the models use different resolutions (128x128 or 64x64) compared to the previous approach (256x256).

Ideas: MSE of cloud coverage is difficult to interpret – a metric normalised per pixel would provide more intuitive values, or Intersection over Union of the cloud coverage masks could be used instead. It would be interesting to also compute FID for the generated conditional image sequences (5.3) – the further images seem blurrier than real data. Using HDR data and adapting the training to handle the dynamic range would be a nice addition (left as future work).

Minor typography notes: use TeX quotes “ and ” (backticks and apostrophes) rather than " (quotes). For a variable sampled from a probability distribution, use $\sim$ (`\sim`) instead of $=$.

Thesis Code

	good	OK	poor	insufficient
Design *... architecture, algorithms, data structures, used technologies*		X		
Implementation *... naming conventions, formatting, comments, testing*		X		
Stability		X		

The algorithm for loading images (section 5.2 in text) does not seem to account for roll-over (9999 → 0000). It would be better to use the newer CSV dataset (`auto_processed`) with timestamps, more images and additional metadata.

Calculating SSIM between consecutive images to detect anomalies is a neat idea.

Documentation: comments, docblocks or Jupyter notebook sections could be used more often. An overview of the classes used during training and inference/sampling, and their tasks would be welcome as well, but is not critical in an experimental thesis where the high-level notebooks are relatively short and self-explanatory.

Overall grade Very Good (better)

Award level thesis No

Date

Signature