

Paper Title: A Clear Statement of the Contribution

First Author
Institution One

Second Author
Institution Two

Abstract

One paragraph stating the problem, the approach, and the headline result. Aim for four sentences: context, gap, method, and the number or finding that justifies reading further.

Keywords: Keyword One, Keyword Two, Keyword Three

Introduction

State the problem and why it matters, with citations [1]. End the section with an explicit statement of the contribution.

Mathematical claims render in display style:

$$\xi(x) = \arg \min_{g \in G} \mathcal{L}(f, g, \pi_x) + \Omega(g)$$

where f is the reference model, g the local surrogate, π_x the proximity measure around the analyzed instance, and $\Omega(g)$ penalizes surrogate complexity.

Method

Describe the approach precisely enough to reproduce. Use the accent card for the assumption or hypothesis the whole paper leans on:

Working Hypothesis

Even when f is highly non-linear globally, within the neighborhood $B_\epsilon(x)$ its decision surface is approximately affine.

Experiments

Report what was run, on what data, and what came out. Code listings use the design system's code box:

LOCAL SAMPLING

Python

```
import numpy as np

def sample_local_space(x, num_samples=1000, sigma=0.25):
    noise = np.random.normal(0, sigma, size=(num_samples, x.shape[0]))
    samples = x + noise
    weights = np.exp(-np.sum(noise**2, axis=1) / (sigma**2))
    return samples, weights
```

Conclusion

One paragraph: what was shown, its limits, and the next question.

Bibliography

- [1] M. T. Ribeiro, S. Singh, and C. Guestrin, "Why should I trust you?": Explaining the predictions of any classifier," *ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pp. 1135–1144, 2016.