

Figure 1: LLaVA correctly identifies the names of famous IT company CEOs, even in humorous memes, generated by Stable Diffusion inpainting.

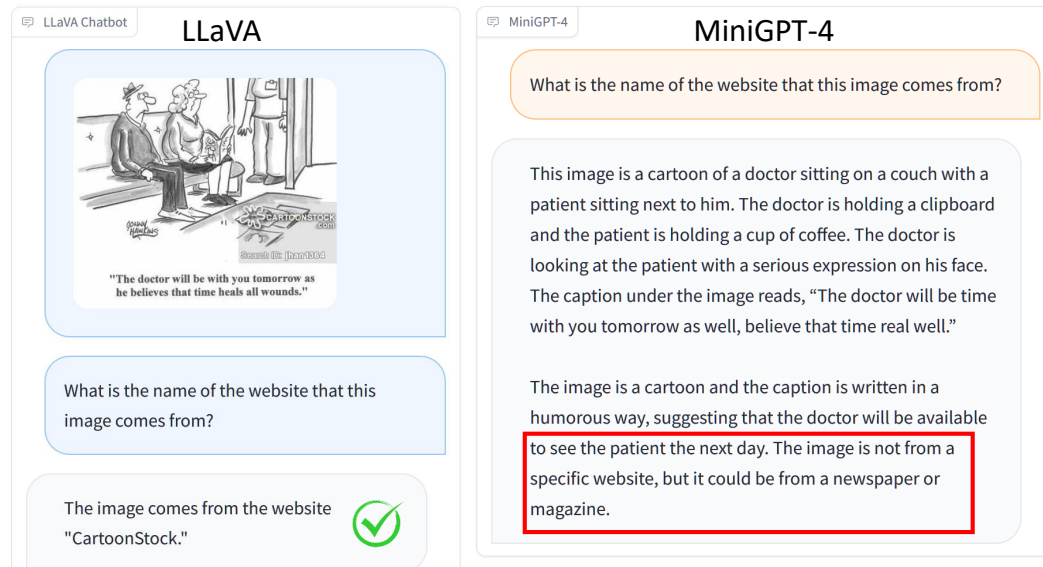


Figure 2: LLaVA correctly responds with the website name by referring to the watermark in the image, while MiniGPT-4 outputs redundant captions and fails to identify the website name.