The Transformer model is the architecture behind many recent breakthroughs, such as LLMs and their various applications. However, in its default definition, it is limited to only one input sequence. At the same time, many real-world phenomena are expressed as multichannel sequences, such as measurements from personal wearable devices or car sensors. This work describes a multichannel Transformer capable of working with such inputs. We evaluate it on the MIMIC-IV dataset on in-hospital mortality and length of stay classification problems. The proposed architecture outperforms baseline models on both problems. We also show that it scales well to sparse inputs, unlike baseline models. Due to the limitations of the dataset and the chosen problems, we cannot make a definitive conclusion about the proposed architecture's general performance on multichannel sequences. We also highlight a need for expert domain knowledge to assess the results.