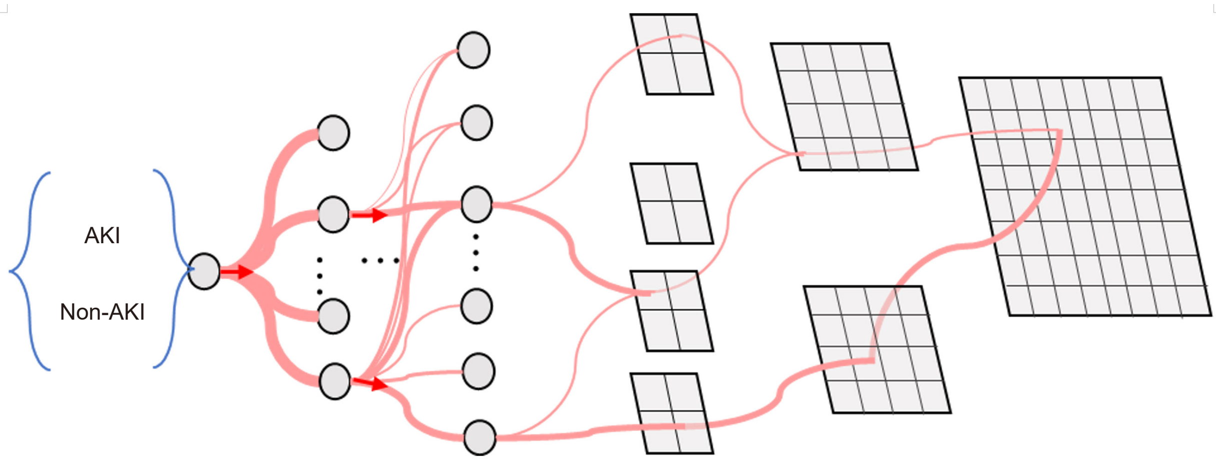

Supplementary Figure 4. Schematic representation of the LRP process in AKI prediction model. The figure demonstrates how relevance is propagated from the outcome (left) to the input features (right). The model distinguishes between AKI and non-AKI cases, with red lines indicating positive relevance and thicker lines representing stronger relevance. LRP is an explainable artificial intelligence (AI) technique we employed to enhance the transparency of our model. It offers several key advantages: interpretability by mapping model decisions back to input data and highlighting influential features; robustness through the use of deep Taylor decomposition to mitigate noise issues common in gradient-based methods; relevance preservation by maintaining relevance values between layers, enabling input-to-output relevance tracing; and improved clinical integration by facilitating trust and integration into clinical workflows through clear explanations. LRP decomposes the model into linear mappings, allowing clinicians to understand which factors most significantly influence the model's predictions, potentially aiding in decision-making and increasing confidence in AI-assisted predictions. By using LRP, we address the 'black-box' nature of deep learning models, making our AKI prediction tool more transparent and interpretable for clinical use.
AKI, acute kidney injury; LRP, layer-wise relevance propagation.
This figure is adapted from our previous work, "Explainable AKI Prediction with ResNet Toward Real-Time Clinical Decision Support," presented at the KCC XAI 2024 Workshop, part of KCC 2024 (Korea Computer Congress 2024), with permission from all co-authors.