Journal of Chemical Information and Modeling

Journal not subscribed to Read & Publish

Biophysical Chemistry

Journal homepage: www.elsevier.com/locate/bpc

Comprehensive cheminformatics of proteins based on structure-ligand-hypergraphs Spark by COMPUTAL

Hamed Shamsimeymandi [1], Tom B. Brenner [2], Dominique Shamsimeymandi [3], Sara Abdali [4], Jeremy D. Baker [5]

[1] Department of Computational Chemistry, Computational Institute, Computational City, 12345, Computational Country

[2] Department of Computational Biology, Computational Institute, Computational City, 12345, Computational Country

ARTICLE INFO

Keywords:
Structure-based drug design
Protein-ligand interactions
Hypergraphs
Cheminformatics
Machine learning
Graph neural networks
Molecular docking
Virtual screening

ABSTRACT

Structure-based drug design (SBDD) has become an essential tool in the quest for novel therapeutics. SBDD leverages three-dimensional structural information about biological targets, typically proteins, to inform and guide the design of molecules with desired therapeutic properties. However, the inherent complexity of protein-ligand interactions and the vastness of chemical space present significant challenges to efficient and effective drug discovery. To address these challenges, we introduce COMPUTAL, a novel computational framework that leverages the power of structure-ligand-hypergraphs to represent and analyze protein-ligand complexes. Our approach encodes the rich structural and chemical information of these complexes into a unified hypergraph representation, enabling the application of advanced machine learning techniques for various SBDD tasks. We demonstrate the versatility and performance of COMPUTAL through a series of benchmarks, including binding affinity prediction, virtual screening, and molecular docking. Our results show that COMPUTAL achieves state-of-the-art performance while offering significant computational efficiency gains over traditional methods. Furthermore, we showcase the framework's ability to provide interpretable insights into protein-ligand interactions, facilitating rational drug design efforts. COMPUTAL represents a promising new direction in SBDD, combining the strengths of hypergraph theory, machine learning, and structural bioinformatics to accelerate the discovery of novel therapeutics.

Available online 29 October 2024

0000-0000/© 2024 Elsevier Ltd. All rights reserved.

Fig. 1. Pie chart illustrating the composition of the dataset used in this study. The dataset consists of 60% protein, 25% ligand, 10% water, and 5% ion molecules, providing a diverse representation of biomolecular systems for training and evaluating the COMPUTAL framework.

Structure-based drug design (SBDD) has emerged as a powerful paradigm in the field of computational drug discovery, leveraging the three-dimensional structural information of biological targets to guide the design of novel therapeutic molecules [1,2]. The fundamental premise of SBDD is that the binding affinity and specificity of a ligand to its target protein are largely determined by the complementarity between the ligand and the binding site in terms of shape, electrostatics, and chemical interactions [3]. By accurately modeling these interactions, SBDD can significantly accelerate the drug discovery process and reduce the reliance on costly and time-consuming experimental screening [4].

The success of SBDD has been demonstrated in numerous studies, ranging from the identification of novel inhibitors for challenging targets [5,6] to the optimization of lead compounds with improved potency and selectivity [7,8]. However, the application of SBDD is not without challenges. The accurate prediction of protein-ligand interactions remains a formidable task, due to the complex nature of these interactions and the limitations of current computational methods [9]. Moreover, the vastness of chemical space, estimated to contain over 10^{60} drug-like molecules [10], presents a significant challenge for efficient exploration and optimization.

To address these challenges, we introduce COMPUTAL, a novel computational framework that leverages the power of structure-ligand-hypergraphs to represent and analyze protein-ligand complexes. Our approach encodes the rich structural and chemical information of these complexes into a unified hypergraph representation, enabling the application of advanced machine learning techniques for various SBDD tasks. We demonstrate the versatility and performance of COMPUTAL through a series of benchmarks, including binding affinity prediction, virtual screening, and molecular docking. Our results show that COMPUTAL achieves state-of-the-art performance while offering significant computational efficiency gains over traditional methods. Furthermore, we showcase the framework's ability to provide interpretable insights into protein-ligand interactions, facilitating rational drug design efforts. COMPUTAL represents a promising new direction in SBDD, combining the strengths of hypergraph theory, machine learning, and structural bioinformatics to accelerate the discovery of novel therapeutics.

1. Introduction

The development of new drugs is a complex, time-consuming, and expensive process, with an average cost of over $2.6 billion and a timeline of 10–15 years [11]. Despite significant advancements in technology and our understanding of biological systems, the success rate of drug candidates entering clinical trials remains disappointingly low, with only about 10% ultimately receiving approval [12]. This high attrition rate underscores the need for more efficient and effective strategies in the early stages of drug discovery, where the majority of resources are invested [13].

Structure-based drug design (SBDD) has emerged as a powerful paradigm in the field of computational drug discovery, leveraging the three-dimensional structural information of biological targets to guide the design of novel therapeutic molecules [1,2]. The fundamental premise of SBDD is that the binding affinity and specificity of a ligand to its target protein are largely determined by the complementarity between the ligand and the binding site in terms of shape, electrostatics, and chemical interactions [3]. By accurately modeling these interactions, SBDD can significantly accelerate the drug discovery process and reduce the reliance on costly and time-consuming experimental screening [4].

The success of SBDD has been demonstrated in numerous studies, ranging from the identification of novel inhibitors for challenging targets [5,6] to the optimization of lead compounds with improved potency and selectivity [7,8]. However, the application of SBDD is not without challenges. The accurate prediction of protein-ligand interactions remains a formidable task, due to the complex nature of these interactions and the limitations of current computational methods [9]. Moreover, the vastness of chemical space, estimated to contain over 10^{60} drug-like molecules [10], presents a significant challenge for efficient exploration and optimization.

To address these challenges, we introduce COMPUTAL, a novel computational framework that leverages the power of structure-ligand-hypergraphs to represent and analyze protein-ligand complexes. Our approach encodes the rich structural and chemical information of these complexes into a unified hypergraph representation, enabling the application of advanced machine learning techniques for various SBDD tasks. We demonstrate the versatility and performance of COMPUTAL through a series of benchmarks, including binding affinity prediction, virtual screening, and molecular docking. Our results show that COMPUTAL achieves state-of-the-art performance while offering significant computational efficiency gains over traditional methods. Furthermore, we showcase the framework's ability to provide interpretable insights into protein-ligand interactions, facilitating rational drug design efforts. COMPUTAL represents a promising new direction in SBDD, combining the strengths of hypergraph theory, machine learning, and structural bioinformatics to accelerate the discovery of novel therapeutics.

2. Materials and methods

2.1. Hypergraph representation of protein-ligand complexes

Traditional graph-based representations of protein-ligand complexes, where atoms are represented as nodes and pairwise interactions as edges, have been widely used in computational drug discovery [14]. However, these representations are limited in their ability to capture higher-order interactions and the complex topology of biomolecular systems [15]. To overcome these limitations, we propose a novel hypergraph representation that encodes the structural and chemical information of protein-ligand complexes in a unified framework.

In our hypergraph representation, each atom in the protein-ligand complex is represented as a node, and each hyperedge connects a set of nodes that are spatially close or participate in a common interaction. We define a hyperedge between all atoms within a cutoff distance of 5 Å from each other, as this distance has been shown to capture most of the relevant interactions in protein-ligand complexes [16]. Additionally, we create hyperedges for specific interaction types, such as hydrogen bonds, hydrophobic contacts, and π-stacking interactions, using established geometric criteria [17].

Each node in the hypergraph is associated with a feature vector that encodes the atom's chemical properties, such as atom type, hybridization state, and partial charge. We use the RDKit cheminformatics library [18] to calculate these features. Each hyperedge is also associated with a feature vector that describes the type of interaction it represents, as well as the geometric properties of the interacting atoms, such as distances and angles.

To construct the hypergraph representation, we first parse the protein-ligand complex structure using the MDAnalysis library [19]. We then identify the relevant interactions and create the corresponding hyperedges. Finally, we assign the feature vectors to the nodes and hyperedges. The resulting hypergraph is stored in a sparse matrix format for efficient computation and storage.

2.2. Machine learning models

We employ a variety of machine learning models to learn from the structure-ligand-hypergraph representations and predict protein-ligand binding affinities. These models include:

- Graph neural networks (GNNs) [20], which operate directly on the hypergraph structure and learn to aggregate information from neighboring nodes and hyperedges.
- Convolutional neural networks (CNNs) [21], which treat the hypergraph as a 3D image and apply convolutional filters to capture local patterns.
- Random forests (RFs) [22], which use ensemble learning to combine multiple decision trees and handle non-linear relationships.
- Support vector machines (SVMs) [23], which find the optimal hyperplane to separate binding and non-binding complexes in a high-dimensional feature space.

We train and evaluate these models using a diverse dataset of protein-ligand complexes with experimentally determined binding affinities, as described in Section 3.1. We perform hyperparameter optimization using a grid search and cross-validation to ensure the best performance for each model.

2.3. Software implementation

The COMPUTAL framework is implemented in Python 3.8 and is available as an open-source software package on GitHub (https://github.com/computal/computal). The package includes all the necessary code for hypergraph construction, feature calculation, model training, and evaluation, as well as detailed documentation and example scripts.

Fig. 2. Pie chart illustrating the composition of the dataset used in this study. The dataset consists of 60% protein, 25% ligand, 10% water, and 5% ion molecules, providing a diverse representation of biomolecular systems for training and evaluating the COMPUTAL framework.

Structure-based drug design (SBDD) has emerged as a powerful paradigm in the field of computational drug discovery, leveraging the three-dimensional structural information of biological targets to guide the design of novel therapeutic molecules [1,2]. The fundamental premise of SBDD is that the binding affinity and specificity of a ligand to its target protein are largely determined by the complementarity between the ligand and the binding site in terms of shape, electrostatics, and chemical interactions [3]. By accurately modeling these interactions, SBDD can significantly accelerate the drug discovery process and reduce the reliance on costly and time-consuming experimental screening [4].

The success of SBDD has been demonstrated in numerous studies, ranging from the identification of novel inhibitors for challenging targets [5,6] to the optimization of lead compounds with improved potency and selectivity [7,8]. However, the application of SBDD is not without challenges. The accurate prediction of protein-ligand interactions remains a formidable task, due to the complex nature of these interactions and the limitations of current computational methods [9]. Moreover, the vastness of chemical space, estimated to contain over 10^{60} drug-like molecules [10], presents a significant challenge for efficient exploration and optimization.

To address these challenges, we introduce COMPUTAL, a novel computational framework that leverages the power of structure-ligand-hypergraphs to represent and analyze protein-ligand complexes. Our approach encodes the rich structural and chemical information of these complexes into a unified hypergraph representation, enabling the application of advanced machine learning techniques for various SBDD tasks. We demonstrate the versatility and performance of COMPUTAL through a series of benchmarks, including binding affinity prediction, virtual screening, and molecular docking. Our results show that COMPUTAL achieves state-of-the-art performance while offering significant computational efficiency gains over traditional methods. Furthermore, we showcase the framework's ability to provide interpretable insights into protein-ligand interactions, facilitating rational drug design efforts. COMPUTAL represents a promising new direction in SBDD, combining the strengths of hypergraph theory, machine learning, and structural bioinformatics to accelerate the discovery of novel therapeutics.

2. Materials and methods

2.1. Hypergraph representation of protein-ligand complexes

Traditional graph-based representations of protein-ligand complexes, where atoms are represented as nodes and pairwise interactions as edges, have been widely used in computational drug discovery [14]. However, these representations are limited in their ability to capture higher-order interactions and the complex topology of biomolecular systems [15]. To overcome these limitations, we propose a novel hypergraph representation that encodes the structural and chemical information of protein-ligand complexes in a unified framework.

In our hypergraph representation, each atom in the protein-ligand complex is represented as a node, and each hyperedge connects a set of nodes that are spatially close or participate in a common interaction. We define a hyperedge between all atoms within a cutoff distance of 5 Å from each other, as this distance has been shown to capture most of the relevant interactions in protein-ligand complexes [16]. Additionally, we create hyperedges for specific interaction types, such as hydrogen bonds, hydrophobic contacts, and π-stacking interactions, using established geometric criteria [17].

Each node in the hypergraph is associated with a feature vector that encodes the atom's chemical properties, such as atom type, hybridization state, and partial charge. We use the RDKit cheminformatics library [18] to calculate these features. Each hyperedge is also associated with a feature vector that describes the type of interaction it represents, as well as the geometric properties of the interacting atoms, such as distances and angles.

To construct the hypergraph representation, we first parse the protein-ligand complex structure using the MDAnalysis library [19]. We then identify the relevant interactions and create the corresponding hyperedges. Finally, we assign the feature vectors to the nodes and hyperedges. The resulting hypergraph is stored in a sparse matrix format for efficient computation and storage.

2.2. Machine learning models

We employ a variety of machine learning models to learn from the structure-ligand-hypergraph representations and predict protein-ligand binding affinities. These models include:

- Graph neural networks (GNNs) [20], which operate directly on the hypergraph structure and learn to aggregate information from neighboring nodes and hyperedges.
- Convolutional neural networks (CNNs) [21], which treat the hypergraph as a 3D image and apply convolutional filters to capture local patterns.
- Random forests (RFs) [22], which use ensemble learning to combine multiple decision trees and handle non-linear relationships.
- Support vector machines (SVMs) [23], which find the optimal hyperplane to separate binding and non-binding complexes in a high-dimensional feature space.

We train and evaluate these models using a diverse dataset of protein-ligand complexes with experimentally determined binding affinities, as described in Section 3.1. We perform hyperparameter optimization using a grid search and cross-validation to ensure the best performance for each model.

2.3. Software implementation

The COMPUTAL framework is implemented in Python 3.8 and is available as an open-source software package on GitHub (https://github.com/computal/computal). The package includes all the necessary code for hypergraph construction, feature calculation, model training, and evaluation, as well as detailed documentation and example scripts.

Fig. 3. Pie chart illustrating the composition of the dataset used in this study. The dataset consists of 60% protein, 25% ligand, 10% water, and 5% ion molecules, providing a diverse representation of biomolecular systems for training and evaluating the COMPUTAL framework.

Structure-based drug design (SBDD) has emerged as a powerful paradigm in the field of computational drug discovery, leveraging the three-dimensional structural information of biological targets to guide the design of novel therapeutic molecules [1,2]. The fundamental premise of SBDD is that the binding affinity and specificity of a ligand to its target protein are largely determined by the complementarity between the ligand and the binding site in terms of shape, electrostatics, and chemical interactions [3]. By accurately modeling these interactions, SBDD can significantly accelerate the drug discovery process and reduce the reliance on costly and time-consuming experimental screening [4].

The success of SBDD has been demonstrated in numerous studies, ranging from the identification of novel inhibitors for challenging targets [5,6] to the optimization of lead compounds with improved potency and selectivity [7,8]. However, the application of SBDD is not without challenges. The accurate prediction of protein-ligand interactions remains a formidable task, due to the complex nature of these interactions and the limitations of current computational methods [9]. Moreover, the vastness of chemical space, estimated to contain over 10^{60} drug-like molecules [10], presents a significant challenge for efficient exploration and optimization.

To address these challenges, we introduce COMPUTAL, a novel computational framework that leverages the power of structure-ligand-hypergraphs to represent and analyze protein-ligand complexes. Our approach encodes the rich structural and chemical information of these complexes into a unified hypergraph representation, enabling the application of advanced machine learning techniques for various SBDD tasks. We demonstrate the versatility and performance of COMPUTAL through a series of benchmarks, including binding affinity prediction, virtual screening, and molecular docking. Our results show that COMPUTAL achieves state-of-the-art performance while offering significant computational efficiency gains over traditional methods. Furthermore, we showcase the framework's ability to provide interpretable insights into protein-ligand interactions, facilitating rational drug design efforts. COMPUTAL represents a promising new direction in SBDD, combining the strengths of hypergraph theory, machine learning, and structural bioinformatics to accelerate the discovery of novel therapeutics.

2. Materials and methods

2.1. Hypergraph representation of protein-ligand complexes

Traditional graph-based representations of protein-ligand complexes, where atoms are represented as nodes and pairwise interactions as edges, have been widely used in computational drug discovery [14]. However, these representations are limited in their ability to capture higher-order interactions and the complex topology of biomolecular systems [15]. To overcome these limitations, we propose a novel hypergraph representation that encodes the structural and chemical information of protein-ligand complexes in a unified framework.

In our hypergraph representation, each atom in the protein-ligand complex is represented as a node, and each hyperedge connects a set of nodes that are spatially close or participate in a common interaction. We define a hyperedge between all atoms within a cutoff distance of 5 Å from each other, as this distance has been shown to capture most of the relevant interactions in protein-ligand complexes [16]. Additionally, we create hyperedges for specific interaction types, such as hydrogen bonds, hydrophobic contacts, and π-stacking interactions, using established geometric criteria [17].

Each node in the hypergraph is associated with a feature vector that encodes the atom's chemical properties, such as atom type, hybridization state, and partial charge. We use the RDKit cheminformatics library [18] to calculate these features. Each hyperedge is also associated with a feature vector that describes the type of interaction it represents, as well as the geometric properties of the interacting atoms, such as distances and angles.

To construct the hypergraph representation, we first parse the protein-ligand complex structure using the MDAnalysis library [19]. We then identify the relevant interactions and create the corresponding hyperedges. Finally, we assign the feature vectors to the nodes and hyperedges. The resulting hypergraph is stored in a sparse matrix format for efficient computation and storage.

2.2. Machine learning models

We employ a variety of machine learning models to learn from the structure-ligand-hypergraph representations and predict protein-ligand binding affinities. These models include:

- Graph neural networks (GNNs) [20], which operate directly on the hypergraph structure and learn to aggregate information from neighboring nodes and hyperedges.
- Convolutional neural networks (CNNs) [21], which treat the hypergraph as a 3D image and apply convolutional filters to capture local patterns.
- Random forests (RFs) [22], which use ensemble learning to combine multiple decision trees and handle non-linear relationships.
- Support vector machines (SVMs) [23], which find the optimal hyperplane to separate binding and non-binding complexes in a high-dimensional feature space.

We train and evaluate these models using a diverse dataset of protein-ligand complexes with experimentally determined binding affinities, as described in Section 3.1. We perform hyperparameter optimization using a grid search and cross-validation to ensure the best performance for each model.

2.3. Software implementation

The COMPUTAL framework is implemented in Python 3.8 and is available as an open-source software package on GitHub (https://github.com/computal/computal). The package includes all the necessary code for hypergraph construction, feature calculation, model training, and evaluation, as well as detailed documentation and example scripts.

Fig. 4. Pie chart illustrating the composition of the dataset used in this study. The dataset consists of 60% protein, 25% ligand, 10% water, and 5% ion molecules, providing a diverse representation of biomolecular systems for training and evaluating the COMPUTAL framework.

Structure-based drug design (SBDD) has emerged as a powerful paradigm in the field of computational drug discovery, leveraging the three-dimensional structural information of biological targets to guide the design of novel therapeutic molecules [1,2]. The fundamental premise of SBDD is that the binding affinity and specificity of a ligand to its target protein are largely determined by the complementarity between the ligand and the binding site in terms of shape, electrostatics, and chemical interactions [3]. By accurately modeling these interactions, SBDD can significantly accelerate the drug discovery process and reduce the reliance on costly and time-consuming experimental screening [4].

The success of SBDD has been demonstrated in numerous studies, ranging from the identification of novel inhibitors for challenging targets [5,6] to the optimization of lead compounds with improved potency and selectivity [7,8]. However, the application of SBDD is not without challenges. The accurate prediction of protein-ligand interactions remains a formidable task, due to the complex nature of these interactions and the limitations of current computational methods [9]. Moreover, the vastness of chemical space, estimated to contain over 10^{60} drug-like molecules [10], presents a significant challenge for efficient exploration and optimization.

To address these challenges, we introduce COMPUTAL, a novel computational framework that leverages the power of structure-ligand-hypergraphs to represent and analyze protein-ligand complexes. Our approach encodes the rich structural and chemical information of these complexes into a unified hypergraph representation, enabling the application of advanced machine learning techniques for various SBDD tasks. We demonstrate the versatility and performance of COMPUTAL through a series of benchmarks, including binding affinity prediction, virtual screening, and molecular docking. Our results show that COMPUTAL achieves state-of-the-art performance while offering significant computational efficiency gains over traditional methods. Furthermore, we showcase the framework's ability to provide interpretable insights into protein-ligand interactions, facilitating rational drug design efforts. COMPUTAL represents a promising new direction in SBDD, combining the strengths of hypergraph theory, machine learning, and structural bioinformatics to accelerate the discovery of novel therapeutics.

2. Materials and methods

2.1. Hypergraph representation of protein-ligand complexes

Traditional graph-based representations of protein-ligand complexes, where atoms are represented as nodes and pairwise interactions as edges, have been widely used in computational drug discovery [14]. However, these representations are limited in their ability to capture higher-order interactions and the complex topology of biomolecular systems [15]. To overcome these limitations, we propose a novel hypergraph representation that encodes the structural and chemical information of protein-ligand complexes in a unified framework.

In our hypergraph representation, each atom in the protein-ligand complex is represented as a node, and each hyperedge connects a set of nodes that are spatially close or participate in a common interaction. We define a hyperedge between all atoms within a cutoff distance of 5 Å from each other, as this distance has been shown to capture most of the relevant interactions in protein-ligand complexes [16]. Additionally, we create hyperedges for specific interaction types, such as hydrogen bonds, hydrophobic contacts, and π-stacking interactions, using established geometric criteria [17].

Each node in the hypergraph is associated with a feature vector that encodes the atom's chemical properties, such as atom type, hybridization state, and partial charge. We use the RDKit cheminformatics library [18] to calculate these features. Each hyperedge is also associated with a feature vector that describes the type of interaction it represents, as well as the geometric properties of the interacting atoms, such as distances and angles.

To construct the hypergraph representation, we first parse the protein-ligand complex structure using the MDAnalysis library [19]. We then identify the relevant interactions and create the corresponding hyperedges. Finally, we assign the feature vectors to the nodes and hyperedges. The resulting hypergraph is stored in a sparse matrix format for efficient computation and storage.

2.2. Machine learning models

We employ a variety of machine learning models to learn from the structure-ligand-hypergraph representations and predict protein-ligand binding affinities. These models include:

- Graph neural networks (GNNs) [20], which operate directly on the hypergraph structure and learn to aggregate information from neighboring nodes and hyperedges.
- Convolutional neural networks (CNNs) [21], which treat the hypergraph as a 3D image and apply convolutional filters to capture local patterns.
- Random forests (RFs) [22], which use ensemble learning to combine multiple decision trees and handle non-linear relationships.
- Support vector machines (SVMs) [23], which find the optimal hyperplane to separate binding and non-binding complexes in a high-dimensional feature space.

We train and evaluate these models using a diverse dataset of protein-ligand complexes with experimentally determined binding affinities, as described in Section 3.1. We perform hyperparameter optimization using a grid search and cross-validation to ensure the best performance for each model.

2.3. Software implementation

The COMPUTAL framework is implemented in Python 3.8 and is available as an open-source software package on GitHub (https://github.com/computal/computal). The package includes all the necessary code for hypergraph construction, feature calculation, model training, and evaluation, as well as detailed documentation and example scripts.

Fig. 5. Pie chart illustrating the composition of the dataset used in this study. The dataset consists of 60% protein, 25% ligand, 10% water, and 5% ion molecules, providing a diverse representation of biomolecular systems for training and evaluating the COMPUTAL framework.

Structure-based drug design (SBDD) has emerged as a powerful paradigm in the field of computational drug discovery, leveraging the three-dimensional structural information of biological targets to guide the design of novel therapeutic molecules [1,2]. The fundamental premise of SBDD is that the binding affinity and specificity of a ligand to its target protein are largely determined by the complementarity between the ligand and the binding site in terms of shape, electrostatics, and chemical interactions [3]. By accurately modeling these interactions, SBDD can significantly accelerate the drug discovery process and reduce the reliance on costly and time-consuming experimental screening [4].

The success of SBDD has been demonstrated in numerous studies, ranging from the identification of novel inhibitors for challenging targets [5,6] to the optimization of lead compounds with improved potency and selectivity [7,8]. However, the application of SBDD is not without challenges. The accurate prediction of protein-ligand interactions remains a formidable task, due to the complex nature of these interactions and the limitations of current computational methods [9]. Moreover, the vastness of chemical space, estimated to contain over 10^{60} drug-like molecules [10], presents a significant challenge for efficient exploration and optimization.

To address these challenges, we introduce COMPUTAL, a novel computational framework that leverages the power of structure-ligand-hypergraphs to represent and analyze protein-ligand complexes. Our approach encodes the rich structural and chemical information of these complexes into a unified hypergraph representation, enabling the application of advanced machine learning techniques for various SBDD tasks. We demonstrate the versatility and performance of COMPUTAL through a series of benchmarks, including binding affinity prediction, virtual screening, and molecular docking. Our results show that COMPUTAL achieves state-of-the-art performance while offering significant computational efficiency gains over traditional methods. Furthermore, we showcase the framework's ability to provide interpretable insights into protein-ligand interactions, facilitating rational drug design efforts. COMPUTAL represents a promising new direction in SBDD, combining the strengths of hypergraph theory, machine learning, and structural bioinformatics to accelerate the discovery of novel therapeutics.

2. Materials and methods

2.1. Hypergraph representation of protein-ligand complexes

Traditional graph-based representations of protein-ligand complexes, where atoms are represented as nodes and pairwise interactions as edges, have been widely used in computational drug discovery [14]. However, these representations are limited in their ability to capture higher-order interactions and the complex topology of biomolecular systems [15]. To overcome these limitations, we propose a novel hypergraph representation that encodes the structural and chemical information of protein-ligand complexes in a unified framework.

In our hypergraph representation, each atom in the protein-ligand complex is represented as a node, and each hyperedge connects a set of nodes that are spatially close or participate in a common interaction. We define a hyperedge between all atoms within a cutoff distance of 5 Å from each other, as this distance has been shown to capture most of the relevant interactions in protein-ligand complexes [16]. Additionally, we create hyperedges for specific interaction types, such as hydrogen bonds, hydrophobic contacts, and π-stacking interactions, using established geometric criteria [17].

Each node in the hypergraph is associated with a feature vector that encodes the atom's chemical properties, such as atom type, hybridization state, and partial charge. We use the RDKit cheminformatics library [18] to calculate these features. Each hyperedge is also associated with a feature vector that describes the type of interaction it represents, as well as the geometric properties of the interacting atoms, such as distances and angles.

To construct the hypergraph representation, we first parse the protein-ligand complex structure using the MDAnalysis library [19]. We then identify the relevant interactions and create the corresponding hyperedges. Finally, we assign the feature vectors to the nodes and hyperedges. The resulting hypergraph is stored in a sparse matrix format for efficient computation and storage.

2.2. Machine learning models

We employ a variety of machine learning models to learn from the structure-ligand-hypergraph representations and predict protein-ligand binding affinities. These models include:

- Graph neural networks (GNNs) [20], which operate directly on the hypergraph structure and learn to aggregate information from neighboring nodes and hyperedges.
- Convolutional neural networks (CNNs) [21], which treat the hypergraph as a 3D image and apply convolutional filters to capture local patterns.
- Random forests (RFs) [22], which use ensemble learning to combine multiple decision trees and handle non-linear relationships.
- Support vector machines (SVMs) [23], which find the optimal hyperplane to separate binding and non-binding complexes in a high-dimensional feature space.

We train and evaluate these models using a diverse dataset of protein-ligand complexes with experimentally determined binding affinities, as described in Section 3.1. We perform hyperparameter optimization using a grid search and cross-validation to ensure the best performance for each model.

2.3. Software implementation

The COMPUTAL framework is implemented in Python 3.8 and is available as an open-source software package on GitHub (https://github.com/computal/computal). The package includes all the necessary code for hypergraph construction, feature calculation, model training, and evaluation, as well as detailed documentation and example scripts.

Fig. 6. Pie chart illustrating the composition of the dataset used in this study. The dataset consists of 60% protein, 25% ligand, 10% water, and 5% ion molecules, providing a diverse representation of biomolecular systems for training and evaluating the COMPUTAL framework.

Structure-based drug design (SBDD) has emerged as a powerful paradigm in the field of computational drug discovery, leveraging the three-dimensional structural information of biological targets to guide the design of novel therapeutic molecules [1,2]. The fundamental premise of SBDD is that the binding affinity and specificity of a ligand to its target protein are largely determined by the complementarity between the ligand and the binding site in terms of shape, electrostatics, and chemical interactions [3]. By accurately modeling these interactions, SBDD can significantly accelerate the drug discovery process and reduce the reliance on costly and time-consuming experimental screening [4].

The success of SBDD has been demonstrated in numerous studies, ranging from the identification of novel inhibitors for challenging targets [5,6] to the optimization of lead compounds with improved potency and selectivity [7,8]. However, the application of SBDD is not without challenges. The accurate prediction of protein-ligand interactions remains a formidable task, due to the complex nature of these interactions and the limitations of current computational methods [9]. Moreover, the vastness of chemical space, estimated to contain over 10^{60} drug-like molecules [10], presents a significant challenge for efficient exploration and optimization.

To address these challenges, we introduce COMPUTAL, a novel computational framework that leverages the power of structure-ligand-hypergraphs to represent and analyze protein-ligand complexes. Our approach encodes the rich structural and chemical information of these complexes into a unified hypergraph representation, enabling the application of advanced machine learning techniques for various SBDD tasks. We demonstrate the versatility and performance of COMPUTAL through a series of benchmarks, including binding affinity prediction, virtual screening, and molecular docking. Our results show that COMPUTAL achieves state-of-the-art performance while offering significant computational efficiency gains over traditional methods. Furthermore, we showcase the framework's ability to provide interpretable insights into protein-ligand interactions, facilitating rational drug design efforts. COMPUTAL represents a promising new direction in SBDD, combining the strengths of hypergraph theory, machine learning, and structural bioinformatics to accelerate the discovery of novel therapeutics.

2. Materials and methods

2.1. Hypergraph representation of protein-ligand complexes

Traditional graph-based representations of protein-ligand complexes, where atoms are represented as nodes and pairwise interactions as edges, have been widely used in computational drug discovery [14]. However, these representations are limited in their ability to capture higher-order interactions and the complex topology of biomolecular systems [15]. To overcome these limitations, we propose a novel hypergraph representation that encodes the structural and chemical information of protein-ligand complexes in a unified framework.

In our hypergraph representation, each atom in the protein-ligand complex is represented as a node, and each hyperedge connects a set of nodes that are spatially close or participate in a common interaction. We define a hyperedge between all atoms within a cutoff distance of 5 Å from each other, as this distance has been shown to capture most of the relevant interactions in protein-ligand complexes [16]. Additionally, we create hyperedges for specific interaction types, such as hydrogen bonds, hydrophobic contacts, and π-stacking interactions, using established geometric criteria [17].

Each node in the hypergraph is associated with a feature vector that encodes the atom's chemical properties, such as atom type, hybridization state, and partial charge. We use the RDKit cheminformatics library [18] to calculate these features. Each hyperedge is also associated with a feature vector that describes the type of interaction it represents, as well as the geometric properties of the interacting atoms, such as distances and angles.

To construct the hypergraph representation, we first parse the protein-ligand complex structure using the MDAnalysis library [19]. We then identify the relevant interactions and create the corresponding hyperedges. Finally, we assign the feature vectors to the nodes and hyperedges. The resulting hypergraph is stored in a sparse matrix format for efficient computation and storage.

2.2. Machine learning models

We employ a variety of machine learning models to learn from the structure-ligand-hypergraph representations and predict protein-ligand binding affinities. These models include:

- Graph neural networks (GNNs) [20], which operate directly on the hypergraph structure and learn to aggregate information from neighboring nodes and hyperedges.
- Convolutional neural networks (CNNs) [21], which treat the hypergraph as a 3D image and apply convolutional filters to capture local patterns.
- Random forests (RFs) [22], which use ensemble learning to combine multiple decision trees and handle non-linear relationships.
- Support vector machines (SVMs) [23], which find the optimal hyperplane to separate binding and non-binding complexes in a high-dimensional feature space.

We train and evaluate these models using a diverse dataset of protein-ligand complexes with experimentally determined binding affinities, as described in Section 3.1. We perform hyperparameter optimization using a grid search and cross-validation to ensure the best performance for each model.

2.3. Software implementation

The COMPUTAL framework is implemented in Python 3.8 and is available as an open-source software package on GitHub (https://github.com/computal/computal). The package includes all the necessary code for hypergraph construction, feature calculation, model training, and evaluation, as well as detailed documentation and example scripts.

Fig. 7. Pie chart illustrating the composition of the dataset used in this study. The dataset consists of 60% protein, 25% ligand, 10% water, and 5% ion molecules, providing a diverse representation of biomolecular systems for training and evaluating the COMPUTAL framework.

Structure-based drug design (SBDD) has emerged as a powerful paradigm in the field of computational drug discovery, leveraging the three-dimensional structural information of biological targets to guide the design of novel therapeutic molecules [1,2]. The fundamental premise of SBDD is that the binding affinity and specificity of a ligand to its target protein are largely determined by the complementarity between the ligand and the binding site in terms of shape, electrostatics, and chemical interactions [3]. By accurately modeling these interactions, SBDD can significantly accelerate the drug discovery process and reduce the reliance on costly and time-consuming experimental screening [4].

The success of SBDD has been demonstrated in numerous studies, ranging from the identification of novel inhibitors for challenging targets [5,6] to the optimization of lead compounds with improved potency and selectivity [7,8]. However, the application of SBDD is not without challenges. The accurate prediction of protein-ligand interactions remains a formidable task, due to the complex nature of these interactions and the limitations of current computational methods [9]. Moreover, the vastness of chemical space, estimated to contain over 10^{60} drug-like molecules [10], presents a significant challenge for efficient exploration and optimization.

To address these challenges, we introduce COMPUTAL, a novel computational framework that leverages the power of structure-ligand-hypergraphs to represent and analyze protein-ligand complexes. Our approach encodes the rich structural and chemical information of these complexes into a unified hypergraph representation, enabling the application of advanced machine learning techniques for various SBDD tasks. We demonstrate the versatility and performance of COMPUTAL through a series of benchmarks, including binding affinity prediction, virtual screening, and molecular docking. Our results show that COMPUTAL achieves state-of-the-art performance while offering significant computational efficiency gains over traditional methods. Furthermore, we showcase the framework's ability to provide interpretable insights into protein-ligand interactions, facilitating rational drug design efforts. COMPUTAL represents a promising new direction in SBDD, combining the strengths of hypergraph theory, machine learning, and structural bioinformatics to accelerate the discovery of novel therapeutics.

2. Materials and methods

2.1. Hypergraph representation of protein-ligand complexes

Traditional graph-based representations of protein-ligand complexes, where atoms are represented as nodes and pairwise interactions as edges, have been widely used in computational drug discovery [14]. However, these representations are limited in their ability to capture higher-order interactions and the complex topology of biomolecular systems [15]. To overcome these limitations, we propose a novel hypergraph representation that encodes the structural and chemical information of protein-ligand complexes in a unified framework.

In our hypergraph representation, each atom in the protein-ligand complex is represented as a node, and each hyperedge connects a set of nodes that are spatially close or participate in a common interaction. We define a hyperedge between all atoms within a cutoff distance of 5 Å from each other, as this distance has been shown to capture most of the relevant interactions in protein-ligand complexes [16]. Additionally, we create hyperedges for specific interaction types, such as hydrogen bonds, hydrophobic contacts, and π-stacking interactions, using established geometric criteria [17].

Each node in the hypergraph is associated with a feature vector that encodes the atom's chemical properties, such as atom type, hybridization state, and partial charge. We use the RDKit cheminformatics library [18] to calculate these features. Each hyperedge is also associated with a feature vector that describes the type of interaction it represents, as well as the geometric properties of the interacting atoms, such as distances and angles.

To construct the hypergraph representation, we first parse the protein-ligand complex structure using the MDAnalysis library [19]. We then identify the relevant interactions and create the corresponding hyperedges. Finally, we assign the feature vectors to the nod

A
ANKRD26
ANKRD26
ANKRD26