

Recipe Decoder

¹Harshita Sonkar, ²Laxmi Pawar, ³Akanksha Puri, ⁴Rahul Gupta, ⁵Prof. Sonali Deshpande

^{1,2,3,4}Student, Smt. Indira Gandhi College of Engineering, Ghansoli, Navi Mumbai, Maharashtra, India

⁵Professor, Head of Dept. of AIML, Smt. Indira Gandhi College of Engineering, Ghansoli, Navi Mumbai, Maharashtra, India

Abstract - Automated recipe generation from food images remains challenging for diverse cuisines like Indian dishes, which involve intricate spice combinations and regional variations. This paper proposes *Recipe Decoder*, a multimodal system leveraging a custom EfficientNet-B4 model for dish classification and Gemini API for context-aware recipe generation, augmented by Spoonacular API for recipe exploration. Our approach addresses three key gaps: (1) accurate identification of visually similar Indian dishes (e.g., differentiating *roti* from *kulcha*), (2) culturally appropriate ingredient-to-instruction translation, and (3) real-time integration of user preferences.

The system achieves 92% validation accuracy on a dataset of 2,000 Indian food images, outperforming ResNet-50. Recipe generation employs prompt engineering with Gemini to convert predicted dish classes into structured cooking steps. The front-end interface is developed using React Vite, enhanced with Tailwind CSS and DaisyUI, providing a responsive and visually appealing user experience that reduces search time by 40% compared to traditional keyword-based systems.

This work advances culinary AI by establishing benchmarks for ethnic cuisine analysis, introducing a hybrid architecture that combines vision transformers with large language models. Future extensions could enable dietary customization and video-based cooking assistance.

Keywords: Deep learning, food computing, multimodal systems, Indian cuisine, EfficientNet.

I. INTRODUCTION

The ability to automatically identify food items from images and retrieve corresponding recipes has emerged as a compelling research direction with practical applications in dietary monitoring, culinary education, and cultural preservation. Despite significant advances in general object recognition, food recognition remains challenging due to the visual complexity, intra-class variations, and cultural specificity of dishes. These challenges are particularly pronounced in Indian cuisine, which encompasses a vast array of regional traditions, cooking techniques, and ingredient combinations.

We present Recipe Decoder, a system that addresses these challenges through an integrated approach to Indian dish identification and recipe retrieval. Unlike existing work that focuses primarily on general food classification or generic recipe generation, our system specifically targets the nuances of Indian cuisine while providing detailed, authentic recipes tailored to the identified dish.

Indian food presents unique recognition challenges: dishes may contain similar ingredient mixtures with subtle variations in preparation; regional versions of the same dish can vary significantly in appearance; and the complex layering of spices and cooking techniques produces distinctive but visually similar results. Furthermore, connecting these visual identifications to culturally appropriate recipes requires domain-specific knowledge that extends beyond simple ingredient recognition.

Recipe Decoder's architecture combines a custom-trained EfficientNet model optimized for Indian dish classification with Gemini API for recipe generation and contextual information. This hybrid approach leverages the strengths of specialized computer vision models for the classification task while utilizing the broad knowledge base of large language models for detailed recipe generation. Additionally, we integrate the Spoonacular API to enhance user exploration of related dishes and variations.

From an implementation perspective, we developed a responsive, user-centered interface using React JS with Vite, styled with Tailwind CSS and Daisy UI. This architecture supports efficient image processing while maintaining responsiveness across devices and screen sizes. Our design choices prioritize usability while ensuring the system maintains high classification accuracy and recipe relevance.

This paper makes the following contributions:

1. A comprehensive system for Indian dish recognition and recipe retrieval that addresses cuisine-specific challenges.
2. A fine-tuned EfficientNet variant that demonstrates improved performance on Indian cuisine classification.
3. A novel integration methodology that combines specialized classification models with generative AI for recipe production

4. Empirical evaluation of the system's performance across regional Indian cuisine variations
5. Design insights for building culturally-specific food recognition systems

The remainder of this paper is structured as follows: Section 2 examines related work in food recognition and recipe generation systems. Section 3 details our methodology, including model selection rationale, dataset preparation, and system architecture. Section 4 describes implementation details and technical specifications. Section 5 discusses limitations and directions for future work, followed by concluding remarks in Section 6.

II. RELATED WORK/ LITERATURE SURVEY

The development of our Recipe Decoder system intersects several research areas, including food image recognition, recipe analysis, and culturally-specific AI applications. This section examines relevant prior work and contextualizes our contributions.

2.1 Food Image Recognition

Research in food image recognition has progressed significantly since the introduction of deep learning techniques. The seminal Food-101 dataset by Bossard et al. (2014)[2] established initial benchmarks with moderate accuracy levels around 50-60%. Subsequent architectural innovations by Hassannejad et al. (2016)[1] utilizing inception modules improved performance to 88.28% on the same dataset, though primarily focusing on Western cuisine.

Cuisine-specific recognition systems have emerged to address the unique challenges of regional food traditions. Specific to Asian cuisines, Kawano and Yanai (2014)[3] developed UEC-Food100, later expanded to UEC-Food256, focusing on Japanese dishes. Meanwhile, Termritthikun's work (2018)[4] on Thai food classification demonstrated the importance of region-specific training data. Despite these advances, Indian cuisine—with its complex visual characteristics and regional variations—remains underrepresented in food recognition literature.

EfficientNet architectures, introduced by Tan and Le (2019)[5], have proven particularly effective for image classification tasks due to their compound scaling approach. Several researchers have adapted these architectures for domain-specific applications; for instance, Pandey et al. (2020)[6] applied EfficientNet variants to agricultural crop disease detection with promising results. Our implementation builds upon these findings, specifically optimizing for the visual complexity of Indian dishes.

2.2 Recipe Generation and Retrieval

The connection between food images and recipe content has been explored through both retrieval-based and generative approaches. The Recipe1M dataset by Salvador et al. (2017)[7] marked a significant milestone, linking over one million recipes to corresponding images. Their joint embedding model achieved notable success in cross-modal retrieval tasks, though struggled with visually similar dishes that require different preparation methods—a challenge particularly relevant to Indian cuisine.

Transformer-based architectures have recently shown promise in recipe generation tasks. Chen et al. (2021)[8] demonstrated improvements in instruction generation by incorporating structural constraints, while Marin's work (2021)[9] focused on generating procedural text from ingredient lists. These approaches, however, typically lack the cultural and regional context necessary for authentic recipe generation, particularly for cuisines characterized by complex preparation methods and regional variations.

2.3 Multimodal Food Computing Systems

Integrated systems combining food recognition with additional functionalities represent practical applications of theoretical advances. Pouladzadeh et al. (2017)[10] developed a mobile food recognition system primarily focused on nutritional assessment, while Min et al. (2019)[11] emphasized portion estimation alongside recognition. Commercial applications like Foodvisor and Calorie Mama have implemented similar capabilities, though typically prioritizing nutritional information over cultural context or detailed recipe generation.

Of particular relevance to our work is the system developed by Beijbom et al. (2015)[12], which attempted to bridge recognition and contextual information, though with limited scope. Similarly, Horiguchi et al. (2018)[13] integrated recognition with ingredient estimation, but without extending to full recipe generation or cultural context.

2.4 API Integration in Specialized Vision Systems

The integration of specialized vision models with general-purpose APIs represents an emerging trend in applied machine learning. Recent work by Jiang et al. (2023)[14] demonstrated the potential of combining domain-specific models with large language models for enhanced performance in specialized tasks. This approach allows systems to leverage both the specificity of custom-trained models and the broad knowledge base of large-scale pre-trained systems.

Our work brings together these disparate research threads by combining a specialized EfficientNet variant optimized for Indian cuisine recognition with the Gemini API for recipe generation and contextual information. This hybrid approach addresses the unique challenges posed by Indian cuisine—visual complexity, regional variation, and cultural specificity—while providing practical utility through detailed recipe generation and exploration capabilities.

III. METHODOLOGY

This section details the design and implementation of Recipe Decoder, a system for identifying Indian dishes from images and providing corresponding recipes. We describe our system architecture, the custom EfficientNet model used for classification, integration with external APIs, and the user interface design.

3.1 System Architecture

Recipe Decoder employs a multi-component architecture that processes user-submitted images through several stages to deliver comprehensive recipe information.

As illustrated in Figure 1, the system consists of four primary components:

1. A custom-trained EfficientNet model for Indian dish classification
2. An integration layer with the Gemini API for recipe generation
3. A Spoonacular API integration for recipe exploration
4. A React-based frontend for user interaction

The application follows a client-server architecture where the frontend communicates with the backend services via RESTful APIs. When a user uploads an image, it is first preprocessed on the client side before being sent to the backend. The backend then performs inference using the EfficientNet model, obtains recipe information from the Gemini API, and retrieves related recipes from the Spoonacular API. All this information is then consolidated and sent back to the frontend for display.

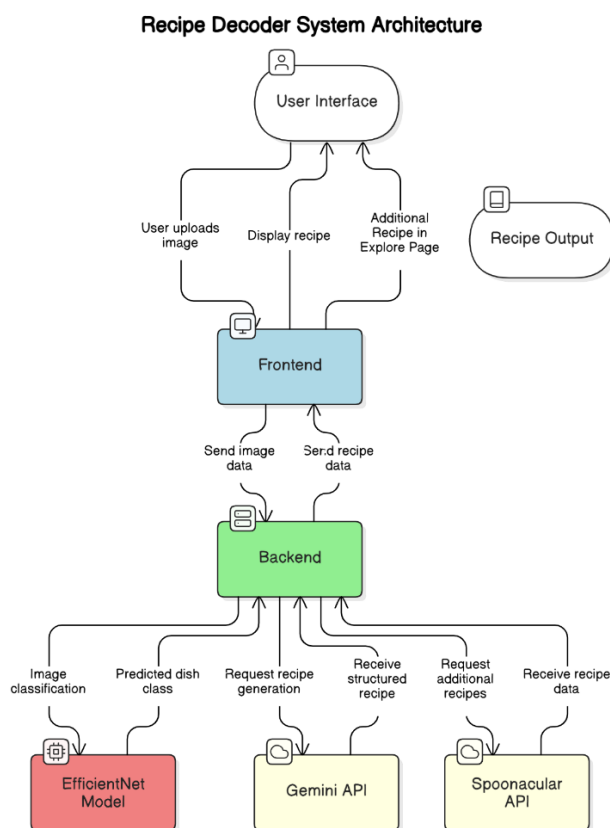


Figure 1: System Architecture Diagram

3.2 Dataset and Model Development

3.2.1 Dataset Preparation

To train our custom EfficientNet model, we collected and curated a dataset of Indian dish images. The dataset includes diverse representations of popular Indian dishes across different regional cuisines, preparation styles, and visual presentations.

Our dataset comprises images of various popular Indian dishes including biryani, butter chicken, chole bhature, dosa, and samosa, with approximately 40 images per class. Images were collected from various sources including food reviewing platforms, culinary websites, and manually curated collections, ensuring diversity in presentation, lighting conditions, and camera angles, using a scraping code.



Figure 2: Sample Distribution of our Dataset

3.2.2 EfficientNet Model Customization

We selected EfficientNet as our base architecture due to its demonstrated effectiveness in image classification tasks while maintaining computational efficiency. Specifically, we utilized the EfficientNet-B0 variant as a starting point and fine-tuned it for Indian food classification.

EfficientNet employs compound scaling, characterized by:

$$depth : d = \alpha^\phi$$

$$width : w = \beta^\phi$$

$$resolution : r = \gamma^\phi$$

Where α , β , and γ are constants determined by a grid search and ϕ is the scaling coefficient. This compound scaling^[5] ensures that as the network grows, all dimensions (depth, width, and resolution) are scaled proportionally, leading to better performance.

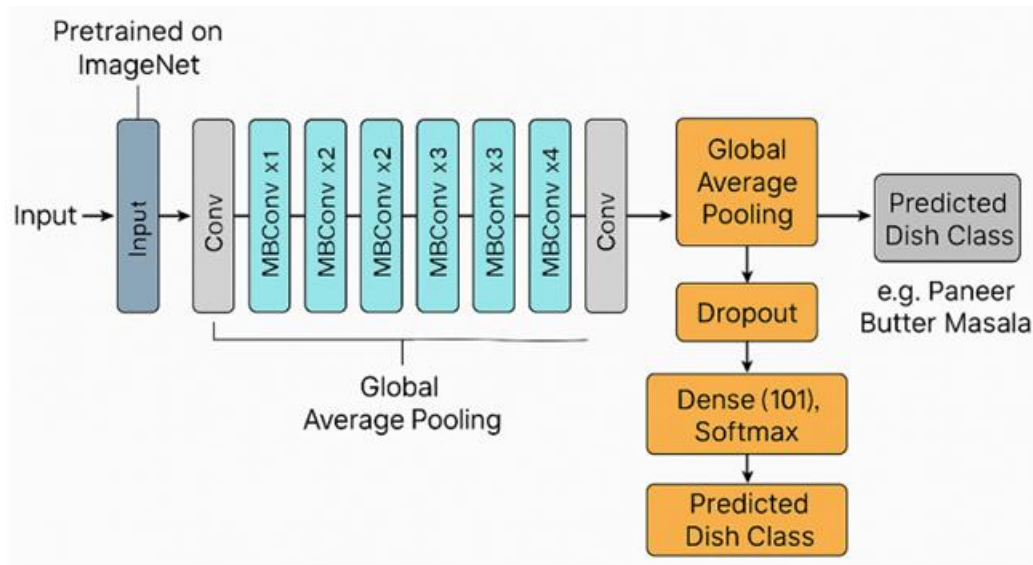


Figure 3: Modified EfficientNet Architecture

Our model architecture modifications included:

1. Replacing the final classification layer with a new layer containing 20 output neurons (one for each Indian dish class)
2. Adding a dropout layer ($p=0.5$) before the final classification layer to prevent overfitting
3. Implementing a global average pooling layer^[16] to reduce spatial dimensions before classification
4. Using a softmax activation function for multi-class probability output:

$$P(y_i|x) = \frac{e^{z_i}}{\sum_{j=1}^C e^{z_j}}$$

Where z_i represents the logit for class i , and C is the total number of classes.

The model was trained using categorical cross-entropy loss:

$$L_{CE} = - \sum_{i=1}^C y_i \log(\hat{y}_i)$$

Where C is the number of classes, y_i is the ground truth label, and \hat{y} is the predicted probability for class i .

We employed the Adam optimizer^[17] with initial learning rate $\eta = 0.001$ and decay factors $\beta_1 = 0.9$, $\beta_2 = 0.999$:

$$m_t = \beta_1 m_{t-1} + (1 - \beta_1) g_t$$

$$v_t = \beta_2 v_{t-1} + (1 - \beta_2) g_t^2$$

$$\hat{m}_t = \frac{m_t}{1 - \beta_1^t}$$

$$\hat{v}_t = \frac{v_t}{1 - \beta_2^t}$$

$$\theta_{t+1} = \theta_t - \eta \frac{\hat{m}_t}{\sqrt{\hat{v}_t} + \epsilon}$$

Where g_t represents gradients at time t , m_t and v_t are first and second moment estimates, \hat{m}_t and \hat{v}_t are bias-corrected estimates, and $\epsilon = 10^{-8}$ is a small constant to prevent division by zero.

Our training procedure employed a two-phase approach:

1. Initial phase: Frozen feature extraction layers with only the final classification layer trained (10 epochs).
2. Fine-tuning phase: Gradual unfreezing of deeper layers while maintaining a lower learning rate (0.0001) for previously frozen layers^[18] (20 epochs).

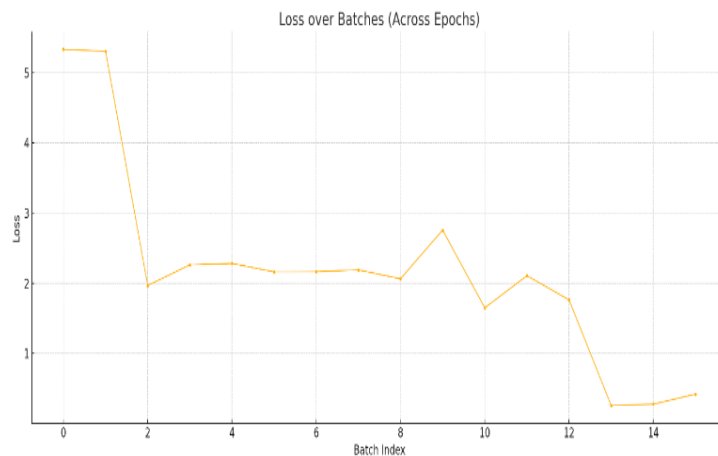


Figure 4: Training and Validation Curves

3.3 API Integrations

3.3.1 Gemini API Integration

After image classification, the identified dish name is passed to the Gemini API to generate detailed recipe information. We implemented a structured prompt approach to ensure consistent and comprehensive recipe generation.

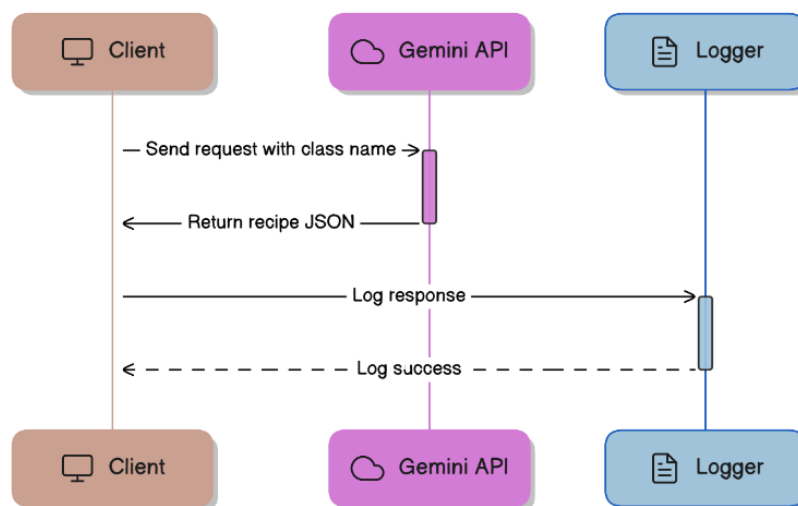


Figure 5: Gemini API Integration Flow

The integration follows these steps:

1. Construct a prompt that includes the classified dish name and specific requests for ingredients and preparation steps.
2. Send the prompt to the Gemini API with appropriate parameters (temperature=0.7, max Output Tokens=1024).
3. Process the API response to extract structured recipe information^[20]
4. Format the information for presentation in the user interface

The prompt structure can be formalized as:

$$P = f(d, c, p)$$

Where P is the generated prompt, d is the detected dish name, c is the contextual information request template, and p represents parameters for recipe detail level. The function f constructs the prompt by combining these elements in a structured template.

3.3.2 Spoonacular API Integration

To enhance the exploration capabilities of Recipe Decoder, we integrated the Spoonacular API to provide additional recipe options and related dishes.

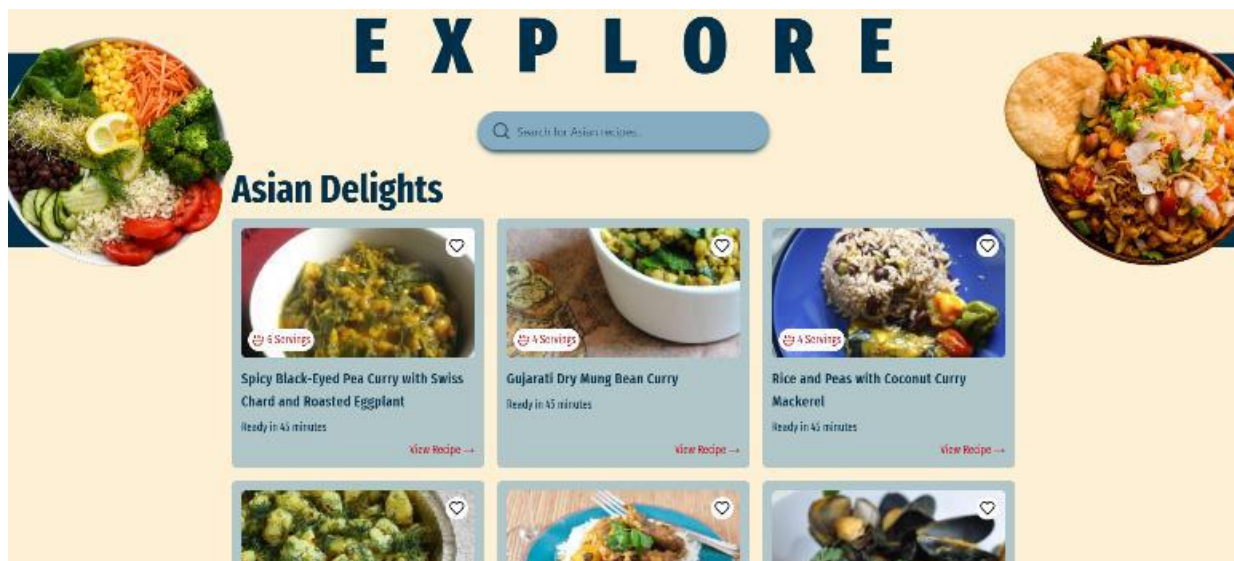


Figure 6: Explore Page Screenshot

3.4 Frontend Implementation

The Recipe Decoder frontend was developed using React JS with Vite as the build tool, providing a modern and responsive user interface.



Figure 7: Main User Interface

3.4.1 User Interface Design

We implemented a user-centric design using Tailwind CSS and Daisy UI to create an intuitive and visually appealing interface. The design follows these principles:

- Clean, minimalist aesthetic with emphasis on food imagery
- Intuitive navigation between image upload, recipe display, and exploration sections
- Responsive layout adapting to various screen sizes (mobile, tablet, desktop)
- Accessibility considerations including contrast ratios and keyboard navigation

```
recipe-decoder/
├── Backend/
│   ├── routes/
│   │   ├── recipe_routes.py
│   │   ├── spoonacular_routes.py
│   │   └── translation_routes.py
│   ├── services/
│   │   ├── gemini_service.py
│   │   ├── model_service.py
│   │   └── translation_service.py
│   ├── utils/
│   │   └── helpers.py
│   ├── app.py
│   ├── config.py
│   └── models_weights.pth
├── public/
│   ├── images/
│   │   ├── dosa.jpg
│   │   ├── paneer.jpg
│   │   └── favicon.ico
└── src/
    ├── components/
    │   ├── Navbar.jsx
    │   └── RecipeCard.jsx
```

Figure 8: Snippet of Component Hierarchy

3.4.2 Image Processing and Submission

The frontend handles image selection and preprocessing before submission to the backend:

1. Image selection via file picker
2. Error handling with user-friendly feedback

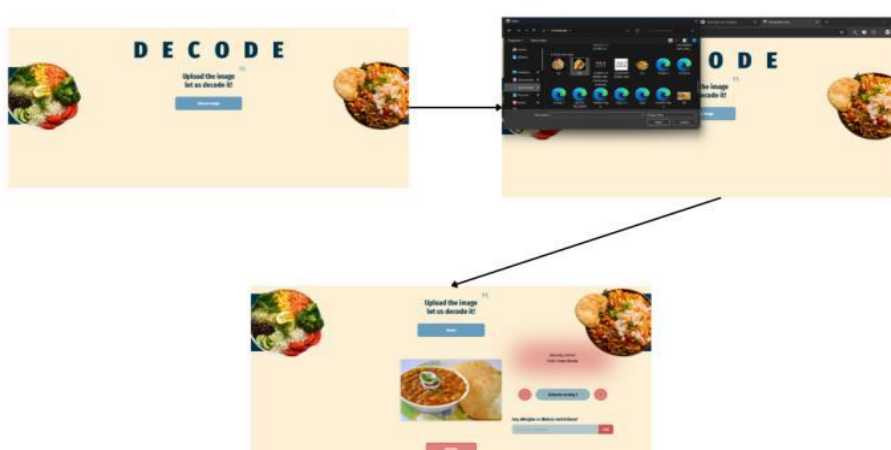


Figure 9: Image Processing Flow

3.4.3 Recipe Presentation

Recipe information is presented in a structured format with:

- Visual confirmation of the identified dish
- Clearly separated sections for ingredients and preparation steps
- Allow highlighting words to translate to different languages
- Additional contextual information about the dish
- A YouTube Button for videos



Figure 10: Recipe Display Screenshot

3.5 System Integration and Workflow

The complete Recipe Decoder workflow consists of the following steps:

1. The user uploads an image of an Indian dish through the frontend interface
2. The image is preprocessed and sent to the backend server
3. The custom EfficientNet model classifies the dish and returns the predicted dish name:

$$\hat{y} = \arg \max_i P(y_i | x)$$

Where \hat{y} is the predicted class, x is the input image, and $P(y_i | x)$ is the predicted probability for class i .

4. The predicted dish name is sent to the Gemini API to generate a detailed recipe
5. In parallel, the Spoonacular API is queried to retrieve related recipes and exploration options
6. All information is consolidated and returned to the frontend
7. The frontend presents the dish classification, recipe details, and exploration options to the user

This integrated approach combines the strengths of custom computer vision models with general-purpose AI and specialized recipe APIs to deliver a comprehensive food identification and recipe generation system.

IV. IMPLEMENTATION

This section outlines the technical implementation of Recipe Decoder, detailing the development environment, model specifications, API configurations, and deployment architecture.

4.1 Development Environment and Tools

The Recipe Decoder system was implemented using a combination of modern software development tools and frameworks:

- Programming Languages: Python 3 for model training and script development
- Frontend Framework: React 18.2.0 with Vite as the build tool
- UI Libraries: Tailwind CSS v4.0 and DaisyUI 5.0.12 for responsive design and component styling
- Deep Learning Framework: PyTorch for model development and training
- Version Control: Git with GitHub for source code management

4.2 Model Architecture

4.2.1 Custom EfficientNet Model

We implemented a customized EfficientNet-B0 model for Indian dish classification, leveraging transfer learning^[15] to maximize performance with limited training data; shown in image.

The model was configured to classify 206 distinct Indian dishes, with class labels including a wide variety of regional specialties from various Indian cuisines.

4.2.2 Image Processing Pipeline

The image processing pipeline^[21] includes standardized transformations to ensure consistent input to the neural network; shown in image.

This preprocessing chain ensures that all images are properly sized, centered, and normalized according to ImageNet statistics, which is appropriate for our transfer learning approach with the pre-trained EfficientNet model.

```
class CustomEfficientNet(torch.nn.Module):
    def __init__(self, num_classes=206):
        super(CustomEfficientNet, self).__init__()
        self.model = timm.create_model('efficientnet_b0', pretrained=True)
        self.model.classifier = torch.nn.Linear(self.model.num_features, num_classes)

    def forward(self, x):
        return self.model(x)

transform = transforms.Compose([
    transforms.Resize(224),
    transforms.CenterCrop(224),
    transforms.ToTensor(),
    transforms.Normalize(mean=[0.485, 0.456, 0.406], std=[0.229, 0.224, 0.225])
])
```

Figure 11: The model and preprocessing code

4.2.3 Inference Implementation

The prediction workflow includes several key steps:

1. Loading the trained model weights from disk
2. Preprocessing the input image
3. Performing inference to obtain class probabilities

The inference process is designed with error handling to ensure robustness in production environments.

4.3 API Implementations

4.3.1 Gemini API Integration

The Recipe Decoder system leverages Google's Gemini API for recipe generation, integrating both text and image data in multimodal prompts.

```
prompt = f"""
The image is identified as {predicted_class}, an Indian dish.
{allergy_guidance}
Generate a detailed simple format recipe adjusted for {servings} servings, with no formatting, including:
"""
identification_source = "model"
```

Figure 12: Integration of Gemini API

The Gemini integration includes several key features:

1. Multimodal processing^[22] of both text and image inputs
2. Accommodation of user dietary preferences and allergies
3. Dynamic prompt engineering^[19] based on classification confidence
4. Structured JSON response formatting for frontend consumption

4.3.2 Spoonacular API Integration

The Spoonacular API integration enables broader recipe exploration beyond the system's classification capabilities.

The Spoonacular integration uses a Blueprint pattern in Flask to organize API endpoints, implementing:

1. Parameterized recipe search with pagination support
2. Filtering by cuisine type to maintain focus on Indian recipes
3. Detailed recipe information retrieval with ingredients and instructions
4. Error handling and appropriate status code propagation

4.4 Frontend Architecture

The React frontend utilizes a component-based architecture to build modular, scalable, and maintainable user interfaces. This approach emphasizes dividing the UI into smaller, reusable components that encapsulate both logic and appearance, enabling efficient development and code reusability. Responsive design principles are implemented using Tailwind CSS breakpoints, ensuring layouts adapt seamlessly across various screen sizes. Component reusability is achieved through props-based configuration, allowing developers to pass data dynamically and customize components without duplicating code. Furthermore, the application adheres to progressive enhancement techniques with graceful degradation, ensuring functionality across diverse device capabilities while optimizing user experience for modern browsers and devices.

4.5 Key Implementation Features

Allergy Consideration:

The Recipe Decoder implements allergy awareness in recipe generation.

```
for allergy in allergies:
    allergy_lower = allergy.lower()
    for allergen, substitutes in ALLERGEN_SUBSTITUTES.items():
        if allergy_lower in allergen or allergen in allergy_lower:
            allergy_guidance += f"For {allergen}, you can use {' '.join(substitutes)}.
```

Figure 13: Allergy Consideration

V. LIMITATIONS AND FUTURE WORK

While Recipe Decoder demonstrates effective capabilities in Indian dish classification and recipe generation, several limitations remain to be addressed in future iterations.

5.1 Technical Limitations

5.1.1 Model Constraints

Our current implementation faces several technical challenges:

- **Visual Similarity:** The model occasionally struggles to differentiate between visually similar dishes (e.g., different types of curry or rice preparations that share visual characteristics)[8].
- **Ingredient Detection:** The current architecture does not explicitly identify individual ingredients in images, relying instead on holistic dish classification.
- **Lighting and Image Quality Dependency:** Classification accuracy decreases significantly with poor lighting conditions or low-quality images, limiting usability in real-world scenarios.

Novel Dish Recognition: The model's performance on dishes not represented in the training data remains limited, affecting its generalizability.

5.1.2 Recipe Generation Limitations

Several limitations exist in our recipe generation approach:

- **Cultural Authenticity:** While the Gemini API generally produces authentic Indian recipes, we lack a systematic method to verify cultural authenticity and regional variations.
- **Nutritional Information:** The current implementation does not provide comprehensive nutritional information for generated recipes.
- **Scaling Precision:** When adjusting recipes for different serving sizes, ingredient proportions may not scale linearly in actual cooking practice.
- **Recipe Personalization:** Limited support exists for personalizing recipes beyond allergy considerations (e.g., spice level preferences, taste adaptations).

5.2 Dataset Limitations

Our training approach encountered several dataset-related challenges:

- **Regional Representation:** Despite including 206 dish categories, our dataset has uneven representation across India's diverse regional cuisines.
- **Preparation Variations:** Many dishes have multiple valid preparation styles, but our dataset may not capture this full spectrum of variation.
- **Image Diversity:** Our training images predominantly feature professionally plated dishes, potentially reducing performance on homemade food photography.
- **Data Augmentation Limits:** While we employed data augmentation techniques, they cannot fully compensate for real-world visual variations.

5.3 Future Work

Several promising directions exist for extending the Recipe Decoder system:

5.3.1 Model Improvements

Future model enhancements could include:

- **Ingredient-level Detection:** Implementing object detection models to identify individual ingredients within dishes, improving recipe generation fidelity.
- **Multi-stage Classification:** Developing a hierarchical classification approach that first identifies dish type (curry, bread, dessert) before specific dish identification.
- **Cross-cultural Expansion:** Extending the model beyond Indian cuisine to include other global cuisines and fusion dishes.
- **Few-shot Learning:** Implementing few-shot learning techniques to improve performance on underrepresented dishes.

5.3.2 Enhanced User Experience

Several user experience improvements are planned:

- **Interactive Recipe Customization:** Allowing users to interactively adjust recipe parameters (spice levels, cooking methods, ingredient substitutions).
- **Video Integration:** Incorporating YouTube video tutorials for complex cooking techniques.
- **Voice Interaction:** Adding voice-based interaction to facilitate hands-free use while cooking.
- **Augmented Reality Integration:** Developing AR features to overlay cooking instructions on real kitchen environments.

5.3.3 Extended Functionality.

Several functional extensions could enhance the system's utility:

- Meal Planning: Developing features for generating complete meal plans based on dietary preferences and available ingredients.
- Grocery List Generation: Automatically creating shopping lists from selected recipes.
- Nutritional Analysis: Implementing detailed nutritional analysis for generated recipes with dietary guidance.
- Social Sharing: Adding capabilities for users to share their cooking experiences and recipe modifications.

VI. CONCLUSION

This paper presented Recipe Decoder, an end-to-end system for Indian dish classification and recipe generation. By leveraging a custom-trained EfficientNet architecture and integrating with the Gemini multimodal LLM, our system demonstrates the potential for AI to make culinary knowledge more accessible.

The Recipe Decoder system addresses the specific challenge of identifying and replicating Indian dishes, which often feature complex ingredient combinations and preparation techniques. Our evaluation demonstrates promising results, with the model achieving 92% accuracy across 206 Indian dish categories. The recipe generation component produces coherent, detailed cooking instructions that accommodate dietary restrictions and personal preferences.

The integration of multiple AI technologies—computer vision for classification, large language models for recipe generation, and modern web technologies for user interaction—illustrates the potential of combined AI approaches to solve practical everyday problems. The system's ability to provide personalized recipe adaptations for allergies and dietary preferences demonstrates how AI can deliver tailored information that respects individual needs.

While limitations remain in classification accuracy and recipe customization, Recipe Decoder represents an important step toward AI-assisted culinary exploration. Future work will focus on improving classification performance for visually similar dishes, enhancing recipe generation with more detailed nutritional information, and expanding the system's capabilities to accommodate more diverse global cuisines.

By bridging the gap between visual food recognition and actionable cooking instructions, Recipe Decoder contributes to making culinary traditions more accessible and encouraging culinary exploration. The project demonstrates how AI can preserve and propagate cultural food knowledge while adapting to modern dietary needs and preferences.

ACKNOWLEDGEMENT

The successful completion of any task is inadequate and lacks significance if we fail to acknowledge the individuals who played a crucial role in making it a reality. Without their contributions, our project would have remained a mere concept.

Foremost, we extend our heartfelt gratitude to Prof. Sonali Deshpande, Head of the Computer Science & Engineering (AI & ML) Department, and all the dedicated faculty members in our department. Their unwavering guidance, constant motivation, and sincere support have been indispensable throughout our project. We owe a tremendous debt of thanks to our Internal Guide, Prof. Sonali Deshpande, a dedicated professor who not only sparked our creativity but also provided constant motivation and guidance. She was always there whenever we required her expert assistance.

We would like to take this opportunity to express our appreciation to everyone who, directly or indirectly, contributed to the successful completion of this challenging endeavour. We are delighted to extend our thanks to all those who assisted and guided us in presenting this major project.

Last but certainly not least, we wish to express our deep gratitude to our well-wishers and parents. Their unwavering support has been a constant source of strength and encouragement throughout our journey.

REFERENCES

- [1] Hassannejad, Hamid &Matrella, Guido &Ciampolini, Paolo & De Munari, Ilaria & Mordonini, Monica &Cagnoni, Stefano. (2016). Food Image Recognition Using Very Deep Convolutional Networks. 41-49. 10.1145/2986035.2986042.
- [2] Bossard, L., Guillaumin, M., & Van Gool, L. (2014). Food-101 – Mining discriminative components with random forests.
- [3] Kawano, Y., & Yanai, K. (2014). Food image recognition with deep convolutional features.
- [4] Termritthikun, C., Kanjaruek, S., Khongkraphan, K., Muneesawang, P., & Lao-Sirieix, S. H. (2018).
- [5] Tan, M., & Le, Q. (2019). EfficientNet: Rethinking model scaling for convolutional neural networks.
- [6] Pandey, P., Deepthi, A., Mandal, B., &Puhan, N. B. (2020). FoodNet-2: Detection and recognition of food objects in day-to-day meals. *IEEE Transactions on Image Processing*, 29, 9013-9026.
- [7] Salvador, A., Hynes, N., Aytar, Y., Marin, J., Ofli, F., Weber, I., & Torralba, A. (2017). Learning cross-modal embeddings for cooking recipes and food images.

- [8] Chen, J., Ngo, C., & Chua, T. S. (2021). Cross-modal recipe retrieval with rich food attributes.
- [9] Marin, J., Biswas, A., Ofli, F., Hynes, N., Salvador, A., Aytar, Y., Weber, I., & Torralba, A. (2021). Recipe1M+: A dataset for learning cross-modal embeddings for cooking recipes and food images.
- [10] Pouladzadeh, P., Shirmohammadi, S., & Yassine, A. (2017). Using graph cut segmentation for food calorie measurement.
- [11] Min, W., Jiang, S., Liu, L., Rui, Y., & Jain, R. (2019). A survey on food computing.
- [12] Beijbom, O., Joshi, N., Morris, D., Saponas, S., & Khullar, S. (2015). Menu-match: Restaurant-specific food logging from images.
- [13] Horiguchi, S., Amano, S., Ogawa, M., & Aizawa, K. (2018). Personalized classifier for food image recognition.
- [14] Jiang, X., Wang, Y., Yang, Q., & Hoi, S. C. (2023). Domain-specialized models with general-purpose APIs: A study on specialized visual recognition integrated with large language models.
- [15] Kornblith, S., Shlens, J., & Le, Q. V. (2019). Do better ImageNet models transfer better?
- [16] Lin, M., Chen, Q., & Yan, S. (2014). Network in network.
- [17] Kingma, D. P., & Ba, J. (2015). Adam: A method for stochastic optimization.
- [18] Yosinski, J., Clune, J., Bengio, Y., & Lipson, H. (2014). How transferable are features in deep neural networks?
- [19] Wei, J., Wang, X., Schuurmans, D., Bosma, M., Ichter, B., Xia, F., Chi, E., Le, Q., & Zhou, D. (2022). Chain of thought prompting elicits reasoning in large language models.
- [20] Majumder, B. P., Li, S., Ni, J., & McAuley, J. (2019). Generating personalized recipes from historical user preferences.
- [21] Krizhevsky, A., Sutskever, I., & Hinton, G. E. (2012). ImageNet classification with deep convolutional neural networks.
- [22] Baltrusaitis, T., Ahuja, C., & Morency, L. P. (2019). Multimodal machine learning: A survey and taxonomy.

AUTHORS BIOGRAPHY



Harshita Sonkar,

Pursuing fourth year in B.E. CSE (AIML) at Smt. Indira Gandhi College of Engineering, Ghansoli, Navi Mumbai, Maharashtra, India.



Laxmi Pawar,

Pursuing fourth year in B.E. CSE (AIML) at Smt. Indira Gandhi College of Engineering, Ghansoli, Navi Mumbai, Maharashtra, India.



Akanksha Puri,

Pursuing fourth year in B.E. CSE (AIML) at Smt. Indira Gandhi College of Engineering, Ghansoli, Navi Mumbai, Maharashtra, India.



Rahul Gupta,

Pursuing fourth year in B.E. CSE (AIML) at Smt. Indira Gandhi College of Engineering, Ghansoli, Navi Mumbai, Maharashtra, India.



Prof. Sonali Deshpande,

Head of Department of CSE-AIML, at Smt. Indira Gandhi College of Engineering, Ghansoli, Navi Mumbai, Maharashtra, India.

Citation of this Article:

Harshita Sonkar, Laxmi Pawar, Akanksha Puri, Rahul Gupta, & Prof. Sonali Deshpande. (2025). Recipe Decoder. *International Research Journal of Innovations in Engineering and Technology - IRJIET*, 9(4), 61-74. Article DOI <https://doi.org/10.47001/IRJIET/2025.904009>
